User Guide for Management Console for Cisco Cloud Email Security (AsyncOS 11.4) - LD (Limited Deployment)

First Published: 2018-03-29

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

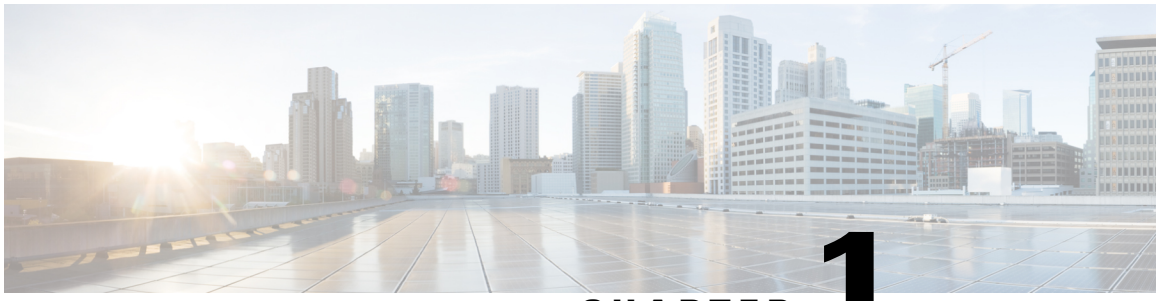

CHAPTER 1

Introduction

This chapter contains the following sections:

What's New in This Release

This section describes the new features and enhancements in this release of AsyncOS for Cloud Email Security Management Console. For more information about the release, see the product release notes, which are available at the following URL:

http://www.cisco.com/c/en/us/support/security/content-security-management-appliance/tsd-products-support-series-home.html

If you are upgrading, you should also review release notes for other releases between your former release and this release, in order to see the features and enhancements that were added in those releases.

Table 1: What's New in this Release

Feature	Description
New Web Interface for Reporting, Quarantine and Tracking	The appliance now has a new web interface to search and view: • Email Reports You can now view email reports from the Reports drop-down based on the following categories: ◦ Email Threat Reports ◦ File and Malware Reports ◦ Connection and Flow Reports ◦ User Reports ◦ Filter Reports For more information, see chapter Using Centralized Email Security Reporting on the Legacy Web Interface, on page 35. • Policy, Virus and Outbreak Quarantines For more information, see chapter Centralized Policy, Virus, and Outbreak Quarantines, on page 197. • Message Tracking For more information, see chapter Tracking Messages, on page 159.

To enable and configure reporting, message tracking, quarantines, network access, and monitor system status, you must access the legacy web interface. Click the gear icon on the Cloud Email Security Management Console, to navigate to the legacy web interface.

The legacy web interface opens in a new browser window and you must log in again to access it.

If you want to log out of the appliance completely, you need to log out of both the new and legacy web interfaces of your appliance.

Changes in Behavior

Change in Report Pages	The following reports are changed in this release: • Overview report page is renamed to Mail Flow Summary. • Outbreak Filters report page is renamed to Outbreak Filtering. • Virus Types report page is renamed to Virus Filtering. • Advanced Malware Protection, AMP File Analysis, AMP Verdict Updates and Mailbox Auto Remediation report pages are merged as Advanced Malware Protection. • Incoming Mail and Outgoing Senders report pages are merged as Mail Flow Details. • TLS Connections report page is renamed to TLS Encryption. • Geo-Distribution report page is renamed to Connection by Country. • Internal Users report page is renamed to User Mail Summary. • Web Interaction Tracking report page is renamed to Web Interaction. For more information, see Understanding the Email Reporting Pages, on page 106.
Changing the User's Password After Expiry	Users are prompted to change the password after the user account is expired. For more information, see Changing the User's Password After Expiry, on page 261.

Comparison of Web Interfaces, AsyncOS 11.4 vs. Previous Releases

The following table shows the comparison of the AsyncOS 11.4 web interface with the previous versions:

Table 2: Comparison of Web Interfaces - AsyncOS 11.4 vs. Previous Releases

Web Interface Page or Element	AsyncOS 11.4	Previous Releases
Landing Page	After you log in to the Cloud Email Security Management Console, the Mail Flow Summary page is displayed.	After you log in to the appliance, the System Status page is displayed.
Management Appliance Tab	Click ⚙ on the Cloud Email Security Management Console to access the Management Appliance tab.	You can enable and configure reporting, message tracking and quarantines, as well as configure network access, and monitor system status.

Web Interface Page or Element	AsyncOS 11.4	Previous Releases
My Reports Page	Click ⚙ on the Cloud Email Security Management Console and choose **Email > Reporting > My Reports** to access the My Reports page.	You can customize your reports dashboard by assembling charts (graphs) and tables from existing report pages.
Reporting Data Availability Page	Click ⚙ on the Cloud Email Security Management Console and choose **Email > Reporting > Reporting Data Availability** to access the Reporting Data Availability page.	You can view, update and sort data to provide real-time visibility into resource utilization and email traffic trouble spots.
Scheduling & Archiving Reports	Click ⚙ on the Cloud Email Security Management Console and choose **Email > Reporting > Scheduled Reports** to schedule your reports. Click ⚙ on the Cloud Email Security Management Console and choose **Email > Reporting > Archive Reports** to archive your reports.	You can schedule reports using the **Email > Reporting > Scheduled Reports** page, and archive your reports using the **Email > Reporting > Archived Report** page of the Security Management appliance.
Reporting Overview Page	The Email Reporting Overview page on the Security Management appliance has been redesigned as Mail Flow Summary page in the new web interface. The Mail Flow Summary page includes trend graphs and summary tables for incoming and outgoing messages.	The Email Reporting Overview page on the Security Management appliance provides a synopsis of the email message activity from your Email Security appliances. The Overview page includes graphs and summary tables for the incoming and outgoing messages.
Report Drill-down	You cannot perform drill-down in reports in the new web interface.	You can perform a drill-down in the following reports: • Incoming Mail • Internal Users • Content Filters • Web Interaction Tracking • System Capacity

Web Interface Page or Element	AsyncOS 11.4	Previous Releases
Advanced Malware Protection Report Pages	The following sections are available on the **Advanced Malware Protection** report page of the Reports menu: • Summary • AMP File Reputation • File Analysis • File Retrospection • Mailbox Auto Remediation	The **Email > Reporting** drop-down menu of the Security Management appliance has the following Advanced Malware Protection report pages: • Advanced Malware Protection • AMP File Analysis • AMP Verdict Updates • Mailbox Auto Remediation
Outbreak Filters Page	The Past Year Virus Outbreaks and Past Year Virus Outbreak Summary are not available in the Outbreak Filtering report page of the new web interface.	The **Email > Reporting Outbreak Filters** page displays the Past Year Virus Outbreaks and Past Year Virus Outbreak Summary.
Spam or End-User Quarantines	To access Spam Quarantine on the Cloud Email Security Management Console, click **Quarantine > Spam Quarantine**. The Spam Quarantine page is displayed in a new browser window.	-
Policy, Virus and Outbreak Quarantines	You can only view Policy, Virus and Outbreak Quarantines in the new web interface. To configure or modify the policy, virus and outbreak quarantines on the Cloud Email Security Management Console, click **Quarantine > Other Quarantine > View** > +.	You can view, configure and modify the Policy, Virus and Outbreak Quarantines on the appliance.
Select All Action for Messages in Quarantine	You can select multiple (or all) messages and perform a message action such as delete, delay, release, move, etc.	You cannot select multiple messages to perform a message action.
Maximum Download Limit for Attachments	The maximum limit for downloading attachments of a quarantined message is restricted to 25 MB.	-

Web Interface Page or Element	AsyncOS 11.4	Previous Releases
Message Details	Click ⚙ on the Cloud Email Security Management Console and choose **Email > Message Quarantine > Policy, Virus and Outbreak Quarantine** to track quarantined messages.	You can perform message tracking of quarantined messages using the Message Details section of the quarantines.
Rejected Connections	To search for rejected connections, click **Tracking > Search > Rejected Connection** tab on the Cloud Email Security Management Console.	-
Query Settings	The Query Settings field of the Message Tracking feature is not available on the Cloud Email Security Management Console.	You can set the query timeout in the Query Settings field of the Message Tracking feature.
Message Tracking Data Availability	Click ⚙ on the Cloud Email Security Management Console and choose **Email > Message Tracking > Message Tracking Data Availability** to access Message Tracking Data Availability page.	You can view the missing-data intervals for your appliance.
Show Additional Details of Messages	You can view additional details of a message such as Verdict Charts, Last State, Sender Groups, Sender IP, SBRS Score and Policy Match details.	-
Verdict Charts and Last State Verdicts	Verdict Chart displays information of the various possible verdicts triggered by each engine in your appliance. Last State of the message determines the final verdict triggered after all the possible verdicts of the engine.	Verdict Charts and Last State Verdicts of the messages are not available.
Message Attachments and Host Names in Message Details	Message attachments and host names are not displayed in the Message Details section of the message on the Cloud Email Security Management Console	Message attachments and host names are displayed in the Message Details section of the message.

Web Interface Page or Element	AsyncOS 11.4	Previous Releases
Sender Groups, Sender IP, SBRS Score and Policy Match in Message Details	Sender Groups, Sender IP, SBRS Score, and Policy Match details of the message is displayed in the Message Details section, on the Cloud Email Security Management Console.	Sender Groups, Sender IP, SBRS Score, and Policy Match of the message is not available in the Message Details section of the message.
Direction of the Message (Incoming or Outgoing)	Direction of the message (incoming or outgoing) is displayed in the message tracking results page, on the Cloud Email Security Management Console.	Direction of the message (incoming or outgoing) is not displayed in the message tracking results page.

Cisco Content Security Management Overview

AsyncOS for Cisco Content Security Management incorporates the following features:

- **External Spam Quarantine:**Hold spam and suspected spam messages for end users, and allow end users and administrators to review messages that are flagged as spam before making a final determination.
- **Centralized Policy, Virus, and Outbreak Quarantines**: Provide a single interface for managing these quarantines and the messages quarantined in them from multiple Email Security appliances. Allows you to store quarantined messages behind the firewall.
- **Centralized reporting:** Run reports on aggregated data from multiple Email Security appliances. The same reporting features available on individual appliances are available on Security Management appliances.
- **Centralized tracking:** Use a single interface to track email messages that were processed by multiple Email Security appliances.
- **Backup of data**: Back up the data on your Security Management appliance, including reporting and tracking data, quarantined messages, and lists of safe and blocked senders.

You can coordinate your security operations from a single Security Management appliance or spread the load across multiple appliances.

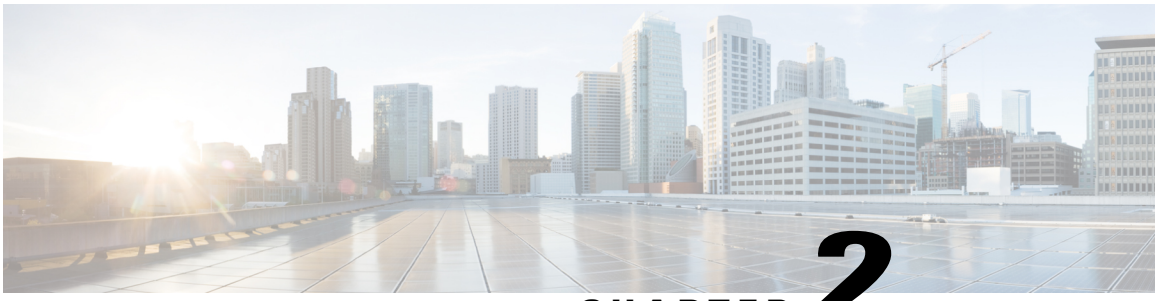

CHAPTER 2

Setup, Installation, and Basic Configuration

This chapter contains the following sections:

Solution Deployment Overview

To configure your Cisco Content Security Management appliance to provide service to your Cisco Content Security solution:

	On These Appliances	Do This	More Information
Step 1	All appliances	Ensure that your appliances meet the system requirements for the features you will use. If necessary, upgrade your appliances.	
Step 2	Email Security appliances	Before you introduce centralized services to your environment, configure all Email Security appliances to provide the security features you want, and verify that all features are working as expected on each appliance.	See the documentation for your Cisco Email Security release.

	On These Appliances	Do This	More Information
Step 3	Security Management appliance	Set up the appliance and run the System Setup Wizard.	See the Installation Planning , on page 10, Preparing for Setup , on page 11 and the Running the System Setup Wizard, on page 16.
Step 4	All appliances	Configure each centralized service that you want to deploy.	Start with the Configuring Services on the Security Management Appliance, on page 20.

Installation Planning

Network Planning

The Security Management appliance lets you separate end user applications from the more secure gateway systems residing in your demilitarized zones (DMZs). Using a two-layer firewall can provide flexibility in network planning so that end users do not connect directly to the outer DMZ .

Figure 1: Typical Network Configuration Incorporating the Security Management appliance

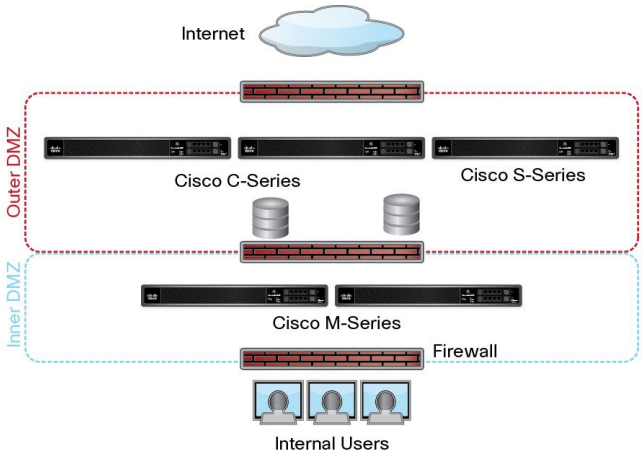

The following figure shows a typical network configuration incorporating the Security Management appliance and multiple DMZs. You deploy the Security Management appliance outside your DMZ, in your internal networks. All connections are initiated by the Security Management appliances (M-Series) to the managed Email Security appliances (C-Series).

Corporate data centers can share a Security Management appliance to perform centralized reporting and message tracking for multiple Email Security appliances. The Security Management appliance can also be used as an external spam quarantine.

After you connect the Email Security appliance to a Security Management appliance and properly configure all appliances, AsyncOS gathers and aggregates data from the managed appliances. From the aggregated data, reports can be generated and an overall view of email usage can be determined.

About Integrating a Security Management Appliance with Email Security Appliances

Additional information about integrating the Security Management appliance with your Email Security appliances, see the "Centralizing Services on a Cisco Content Security Management Appliance" chapter in the user documentation or online help for your Email Security appliance.

Deployments with Clustered Email Security Appliances

The Security Management appliance cannot be placed in a cluster of Email Security appliances that are using the email appliances' centralized management feature. However, clustered Email Security appliances can deliver messages to the Security Management appliance for centralized reporting and tracking and to quarantine messages.

Preparing for Setup

Before you run the System Setup Wizard:

Step 1 Review the latest release notes for your product. See Network Planning , on page 10.

Step 2 Verify that the components of your security solution are compatible. .

Step 3 Ensure that your network and physical space are ready to support this deployment. See Installation Planning , on page 10.

Step 4 Physically set up and connect the Security Management appliance. See Physically Setting Up and Connecting the Appliance , on page 11.

Step 5 Determine network and IP address assignments. See Determining Network and IP Address Assignments, on page 12.

Step 6 Gather information about your system setup. See Gathering the Setup Information, on page 12.

Physically Setting Up and Connecting the Appliance

Before you follow the procedures in this chapter, complete the steps described in the quick start guide that came with your appliance. In this guide, it is assumed that you have unpacked the appliance, physically installed it in a rack, and turned it on.

Before you can log in to the GUI, you need to set up a private connection between a PC and the Security Management appliance. For example, you can use the included crossover cable to connect directly from the Management port on the appliance to a laptop. Optionally, you can connect through an Ethernet connection between a PC and the network (for example, an Ethernet hub) and between the network and the Management port on the Security Management appliance.

Determining Network and IP Address Assignments

Note If you have already cabled your appliance to your network, ensure that the default IP address for the content security appliance does not conflict with other IP addresses on your network. The IP address that is pre-configured on the Management port of each appliance is 192.168.42.42.

After setup, go to the **Management Appliance > Network > IP Interfaces** page on the main Security Management appliance to change the interface that the Security Management appliance uses.

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

- IP address
- Netmask

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

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

For more information, see Assigning Network and IP Addresses, on page 397.

Note If you are running a firewall on your network between the Internet and the content security appliance, it may be necessary to open specific ports for the appliance to work properly. For more information on firewalls, see Firewall Information, on page 401

Always use the same IP address on the Security Management appliance for receiving and sending email messages to the Email Security appliances. For an explanation, see information about Mail Flow in the documentation for your Email Security appliance.

Note that IPv6 is not supported for communication between the Cisco Content Security Management appliance and the appliances it manages.

Gathering the Setup Information

Use the following table to gather information about system setup. You will need this information at hand while running the System Setup Wizard.

Note See the Assigning Network and IP Addresses, on page 397 for detailed information about network and IP addresses.

The following table shows the system setup worksheet

1	Notifications		Email address where system alerts are sent:

2	System Time		NTP Server (IP address or hostname):
3	Admin Password		Choose a new password for the “admin” account:
4	AutoSupport		Enable AutoSupport? ___ Yes ___ No
5	Hostname		Fully qualified hostname of the Security Management appliance:
6	Interface / IP Address		IP address:
			Netmask:
7	Network	Gateway	Default Gateway (router) IP address:
		DNS	___ Use the Internet’s root DNS servers
			___ Use these DNS servers:

Accessing the Security Management Appliance

The Security Management appliance has a standard web-based graphical user interface, a separate web-based interface for managing the spam quarantine, a command-line interface, and special or limited web interfaces for administrative users granted access to specific features and functionality.

Browser Requirements

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

Table 3: Supported Browsers and Releases

Browser	Windows 7	MacOS 10.6
Safari	—	7.0 and later
Google Chrome	Latest Stable Version	Latest Stable Version
Microsoft Internet Explorer	11.0	—

Browser	Windows 7	MacOS 10.6
Mozilla Firefox	Latest Stable Version	Latest Stable Version

- Internet Explorer 11.0 (Windows 7 only)
- Safari (7 and later)
- Firefox (Latest Stable Version)
- Google Chrome (Latest Stable Version)

Browsers are supported only for operating systems officially supported by the browser.

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

Note The best viewed resolution for the AsyncOS (11.4 and above) web interface is 1366x786, for all the browsers.

About Accessing the Web Interfaces

The Security Management appliance has two web interfaces: the standard administrator interface, available by default on port 80, and the spam quarantine end user interface, available by default on port 82. The spam quarantine HTTPS interface defaults to port 83 when enabled.

Because you can specify HTTP or HTTPS when configuring each of the web interfaces (go to **Management Appliance > Network > IP Interfaces** on the Security Management appliance), you may be asked to reauthenticate if you switch between the two during your session. For example, if you access the admin web interface through HTTP on port 80 and then, in the same browser, access the spam quarantine end user web interface through HTTPS on port 83, you are asked to reauthenticate if you return to the admin web interface.

Note - When accessing the GUI, do not use multiple browser windows or tabs simultaneously to make changes to the Security Management appliance. Do not use concurrent GUI and CLI sessions either. Doing so will cause unexpected behavior and is not supported.

- By default, your session times out if you are idle for more than 30 minutes or if you close the browser without logging out. If this happens, you must reenter your user name and password. To change the timeout limit, see Configuring the Web UI Session Timeout, on page 276.

Accessing the Web Interface

Step 1 Open your web browser and type `192.168.42.42` in the IP address text field.

Step 2 Enter the following default values:

- User name: `admin`
- Password: `ironport`

Note This password is NOT valid after you complete the System Setup Wizard, either using the web interface or the command-line interface.

Accessing the Command Line Interface

The command line interface, or CLI, is accessed on the Security Management appliance in the same way that the CLI is accessed on all Cisco Content Security appliances. There are, however, some differences:

- System setup must be performed through the GUI.
- Some CLI commands are not available on the Security Management appliance. For a list of which commands are not supported, see the IronPort AsyncOS CLI Reference Guide for Cisco Content Security Appliances.

For production deployments, you should use SSH to access the CLI. Use a standard SSH client to access the appliance on port 22. For lab deployments, you can also use telnet; however, this protocol is not encrypted.

Supported Languages

With the appropriate license key, AsyncOS can display the GUI and CLI in any of the following languages:

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

To choose the GUI and default reporting language, do one of the following:

- Set the language preference. See Setting Preferences , on page 346.
- Use the Options menu at the top right side of the GUI window to select the language for the session.

(The method that works depends on the method used to authenticate your login credentials.)

Running the System Setup Wizard

AsyncOS provides a browser-based System Setup Wizard to guide you through the process of system configuration. Later, you may want to take advantage of custom configuration options not available in the wizard. However, you must use the wizard for the initial setup to ensure a complete configuration.

The Security Management appliance supports this wizard via the GUI only. It does not support system setup through the command line interface (CLI).

Before You Begin

Complete all tasks in the Preparing for Setup , on page 11.

Caution The System Setup Wizard completely reconfigures the appliance. Only use the wizard when you initially install the appliance, or if you want to completely overwrite the existing configuration.

Be sure to connect the Security Management appliance to your network through the Management port.

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

Note By default, your session times out if you are idle for more than 30 minutes or if you close the browser without logging out. If this happens, you must reenter your user name and password. If the session times out while you are running the System Setup Wizard, you need to start over from the beginning. To change the timeout limit, see Configuring the Web UI Session Timeout, on page 276.

Overview of the System Setup Wizard

Step 1 Launch the System Setup Wizard , on page 17

Step 2 Review the End User License Agreement, on page 17

Step 3 Configure the System Settings, on page 17

- Notification settings and AutoSupport
- System time settings
- Admin password

Step 4 Configure the Network Settings, on page 18

- Hostname of the appliance
- IP address, network mask, and gateway of the appliance

- Default router and DNS settings

Step 5 Review Your Configuration, on page 18
Proceed through the wizard pages, and carefully review your configuration at step 4. You can return to a step by clicking **Previous**. At the end of the process, the wizard prompts you to commit the changes that you have made. Most changes do not take effect until you commit them.

Step 6 Proceeding to the Next Steps, on page 19

Launch the System Setup Wizard

To launch the wizard, log in to the GUI as described in the Accessing the Web Interface , on page 14. The first time you log in to the GUI, the initial page of the System Setup Wizard appears by default. You can also access the System Setup Wizard from the System Administration menu (Management Appliance > System Administration > System Setup Wizard).

Review the End User License Agreement

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

Configure the System Settings

Entering an Email Address for System alertsAlerts

AsyncOS sends alert messages through email if there is a system error that requires your intervention. Enter the email address (or addresses) where the alerts are sent.

You need to add at least one email address for the system alerts. Separate multiple addresses with commas. The email addresses that you enter initially receive all types of alerts at all levels. You can customize the alert configuration later. For more information, see the Managing Alerts, on page 312.

Setting the Time

Set the time zone on the Security Management appliance so that timestamps in reports, message headers and log files are correct. Use the drop-down menus to locate your time zone or to define the time zone by GMT offset.

You can set the system clock time manually, but Cisco recommends using an Network Time Protocol (NTP) server to synchronize time with other servers on your network or the Internet. By default, the Cisco NTP server (time.sco.cisco.com) is added as an entry to synchronize the time on your content security appliance. Enter the hostname of the NTP server, and click Add Entry to configure an additional NTP server. For more information, see the Configuring the System Time, on page 324.

Setting the Password

You must change the password:adminpassword for the AsyncOS admin account. Keep the password in a secure location. Changes to the password take effect immediately.

Note If you cancel the system setup after resetting the password, your password changes are not undone.

Enabling AutoSupport

The AutoSupport feature (enabled by default) notifies Customer Support about issues with the Security Management appliance so that they can provide optimal support. For more information, see the Cisco AutoSupport, on page 315.

Configure the Network Settings

Define the hostname of the machine and then configure the gateway and DNS settings.

Note Verify that you have connected the Security Management appliance to your network through the Management port.

Network Settings

Enter the fully qualified hostname for the Security Management appliance. This name should be assigned by the network administrator.

Enter the IP address of the Security Management appliance.

Enter the network mask and IP address of the default deerrouter (gateway) on your network.

Next, configure the Domain Name Service (DNS) settings. AsyncOS contains a high-performance internal DNS resolver/cache that can query the Internet's root servers directly, or the system can use DNS servers that you specify. If you use your own servers, you need to supply the IP address of each DNS server. You can enter up to four DNS servers when you are using the System Setup Wizard.

Note The DNS servers you specify have an initial priority of 0. For more information, see the Configuring Domain Name System Settings, on page 320.

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

Review Your Configuration

Now, the System Setup Wizard displays a summary of the setup information that you have entered. If you need to make any changes, click **Previous** at the bottom of the page and edit the information.

After you have reviewed the information, click **Install This Configuration**. Then click **Install** in the confirmation dialog box that appears.

If the page appears not to respond when you click **Install This Configuration**, this is because the appliance is now using the the new IP address that you specified in the wizard. To continue using the appliance, use the

new IP address. If you followed the instructions in the Quick Start Guide to temporarily change the IP address of the computer you used to access your new hardware appliance, revert your computer's IP address to its original settings first.

Proceeding to the Next Steps

After you install the Security Management appliance and run the System Setup Wizard, you can modify other settings on the appliance and configure the monitoring services.

Depending on the process you used to access the appliance in order to run the system setup wizard, the **System Setup Next Steps** page appears. If this page dos not appear automatically, you can access it by choosing **Management Appliance > System Administration > Next Steps**.

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

About Adding Managed Appliances

You will add managed Email appliances to the Security Management appliance when you configure the first centralized service for each appliance.

When you add a remote appliance, the Security Management appliance compares the product name of the remote appliance with the type of appliance you are adding. The Security Management appliance will also check the monitoring services on the remote appliances to make sure that they are correctly configured and compatible.

The Security Appliances page shows the managed appliances that you have added. The Connection Established? column shows whether or not the connection for monitoring services is properly configured.

Instructions for adding managed appliances are included in the following procedures:

- Adding the Centralized Email Reporting Service to Each Managed Email Security Appliance , on page 103
- Adding the Centralized Message Tracking Service to Each Managed Email Security Appliance , on page 162
- Adding the Centralized Spam Quarantine Service to Each Managed Email Security Appliance , on page 177
- Adding the Centralized Policy, Virus, and Outbreak Quarantine Service to Each Managed Email Security Appliance, on page 202

Editing Managed Appliance Configurations

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Security Appliances**.

Step 3 In the Security Appliance section, click on the name of the appliance you want to edit.

Step 4 Make the necessary changes to the appliance configuration.

For example, select or clear check boxes for monitoring services, reconfigure file transfer access, or change the IP address.

Note Changing the IP address of a managed appliance can cause several issues to occur. If you change the IP address of an Email Security appliance, the tracking availability data for the appliance will be lost.

Step 5 Click **Submit** to submit your changes on the page, then click Commit Changes to commit your changes.

Removing an Appliance from the List of Managed Appliances

Before You Begin

You may need to disable any enabled centralized services on the remote appliance before you can remove that appliance from the Security Management appliance. For example, if the Centralized Policy, Virus, and Outbreak Quarantine service is enabled, you must disable that service first on the Email Security appliance. See the documentation for your email appliance.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Security Appliances**.

Step 3 In the Security Appliances section, and click the trash can icon in the row for the managed appliance that you want to delete.

Step 4 In the confirmation dialog box, click **Delete.**

Step 5 Submit and commit your changes.

Configuring Services on the Security Management Appliance

Email security services:

- Using Centralized Email Security Reporting on the Legacy Web Interface, on page 35
- Tracking Messages, on page 159
- Spam Quarantine, on page 173
- Centralized Policy, Virus, and Outbreak Quarantines, on page 197

Committing and Abandoning Configuration Changes

After you make most configuration changes in the Cisco Content Security appliance GUI, you must explicitly commit the changes.

Figure 2: The Commit Changes Button

To	Do This
Commit all pending changes	Click the orange Commit Changes button at the top right side of the window. Add a description of the changes and then click commit. If you have not made any changes that require a commit, then a gray No Changes Pending button appears instead of Commit Changes.
Abandon all pending changes	Click the orange Commit Changes button at the top right side of the window, then click Abandon Changes.

Note The configuration changes made on the old web interface is updated on the new web interface, after you logout and login to the new Cisco Content Security Management web interface.

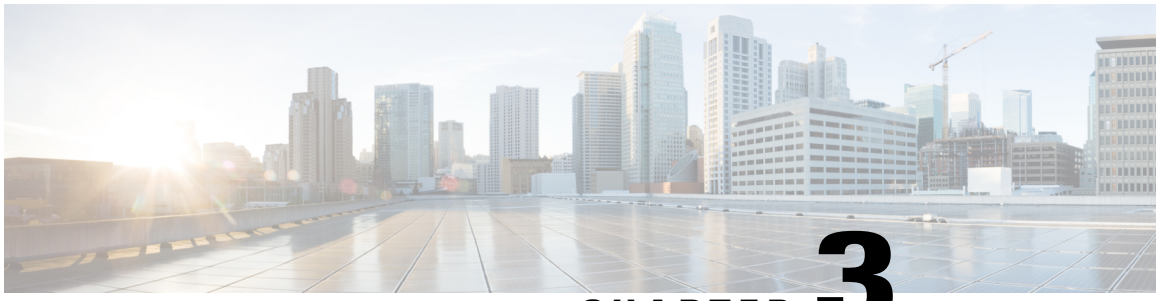

CHAPTER 3

Working With Reports on the Legacy Web Interface

This chapter contains the following sections:

Ways to View Reporting Data

Table 4: Ways To View Reporting Data

To	See
View and customize web-based interactive report pages	• Customizing Your View of Report Data , on page 25 • Using Centralized Email Security Reporting on the Legacy Web Interface, on page 35 • Centralized Policy, Virus, and Outbreak Quarantines, on page 197

To	See
Automatically generate recurring PDF or CSV reports	• Scheduling Email Reports, on page 83
Generate a PDF or CSV report on demand	• Generating Email Reports On Demand , on page 85
Export raw data as a CSV (Comma-separated values) file	• Printing and Exporting Reporting Data , on page 31 • Exporting Report Data as a Comma Separated Values (CSV) File , on page 32
Generate a PDF of report data	Printing and Exporting Reporting Data , on page 31
Email report information to yourself and other people	• Generating Email Reports On Demand , on page 85 • Scheduling Email Reports, on page 83
Find information about specific transactions	• Viewing Details of Messages or Transactions Included in Reports , on page 30

Note For differences between logging and reporting, see Logging Versus Reporting, on page 347.

How the Security Management Appliance Gathers Data for Reports

The Security Management appliance pulls data for all reports from all managed appliances approximately every 15 minutes and aggregates the data from these appliances. Depending on your appliance, it may take awhile for a particular message to be included in the reporting data on the Security Management appliance. Check the **System Status** page for information on your data.

Reporting Data includes transactions involving both IPv4 and IPv6.

Note When gathering data for reports, the Security Management appliance applies the timestamp from the information that was set when you configured the time settings on the Security Management appliance. For information on setting the time on your Security Management appliance, see the Configuring the System Time, on page 324.

How Reporting Data is Stored

All of the appliances store reporting data. The following table shows what time periods that each appliance stores data.

Table 5: Reporting Data Storage on the Email and Web Security Appliances

	Minute	Hourly	Daily	Weekly	Monthly	Yearly
Local Reporting on Email Security appliance **or** Web Security appliance	•	•	•	•	•	
Centralized Reporting on Email Security appliance **or** Web Security appliance	•	•	•	•		
Security Management appliance		•	•	•	•	•

About Reporting and Upgrades

New reporting features may not apply to transactions that occurred before upgrade, because the required data may not have been retained for those transactions. For possible limitations related to reporting data and upgrades, see the Release Notes for your release.

Customizing Your View of Report Data

When viewing report data in the web interface, you can customize your view.

To	Do This
View data per appliance or reporting group	Viewing Reporting Data for an Appliance or Reporting Group, on page 26
Specify a time range	Choosing a Time Range for Reports , on page 26
(For Web reports) Choose which data to chart	(Web Reports Only) Choosing Which Data to Chart , on page 27

To	Do This
Customize tables	See Customizing Tables on Report Pages , on page 27
Search for specific information or a subset of data to view	• For Email reports, Searching and the Interactive Email Report Pages , on page 39. • For Web reports, look for a Find or Filter option at the bottom of most tables. • Some tables include links (in blue text) to details for aggregated data.
Specify report-related preferences	See Setting Preferences , on page 346
Create a custom report with only the charts and tables you want	See Custom Reports , on page 28.

Note Not all customization features are available for every report.

Viewing Reporting Data for an Appliance or Reporting Group

For Mail Flow Summary and System Capacity reports for Email, you can view data from all appliances, or from any one centrally-managed appliance.

For Email reports, if you have created groups of Email Security appliances as described in Creating Email Reporting Groups, on page 37, you can view the data for each reporting group.

To specify the view, select an appliance or group from the **View Data For** list on supported pages.

If you are viewing report data on the Cloud Email Security Management Console to which you have recently taken backup from another Security Management appliance, you must first add (but do not establish a connection to) each appliance in > **Management Appliance** > **Centralized Services** > **Security Appliances**.

Choosing a Time Range for Reports

Most predefined report pages allow you to choose a Time Range for the data to include. The time range that you select is used for all of the report pages until you select a different value in the Time Range menu.

Available Time Range options differ by appliance and differ for Email reporting on the Security Management appliance:

Note Time ranges on report pages are displayed as a Greenwich Mean Time (GMT) offset. For example, Pacific time is GMT + 7 hours (GMT + 07:00).

Note All reports display date and time information based on the systems configured time zone, shown as a Greenwich Mean Time (GMT) offset. However, data exports display the time in GMT to accommodate multiple systems in multiple time zones around the world.

Tip You can specify a default time range that will always display each time you log in. For information, see Setting Preferences , on page 346.

(Web Reports Only) Choosing Which Data to Chart

The default charts on each Web Reporting page display commonly-referenced data, but you can choose to chart different data instead. If a page has multiple charts, you can change each chart.

Generally, the chart options are the same as the columns of the table in the report. However, some columns cannot be charted.

Charts reflect all available data in a table column, regardless of the number of items (rows) you choose to display in the associated table.

Step 1 Click the **Chart Options** link below a chart.

Step 2 Choose the data to display.

Step 3 Click **Done**.

Customizing Tables on Report Pages

Table 6: Customizing Tables on Web Report Pages

To	Do This	More Information
• Show additional columns • Hide visible columns • Determine available columns for a table	Click the **Columns** link below the table, select the columns to display, then click **Done**.	For most tables, some columns are hidden by default. Each each report page offers different columns. See also Table Column Descriptions for Email Reporting Pages, on page 43.
Reorder table columns	Drag a column heading to the desired new position	—
Sort the table by the heading of your choice.	Click a column heading.	—

To	Do This	More Information
Display more or fewer rows of data	From the **Items Displayed** drop-down list at the top right of a table, choose a number of rows to display.	For Web reports, you can also set a preference for a default number of rows to display; see Setting Preferences , on page 346.
View details about a table entry, where available	Click a blue entry in the table	See also Viewing Details of Messages or Transactions Included in Reports , on page 30.
Narrow the pool of data to a specific subset	Choose or enter a value in the filter setting below the table, where available	For Web reports, available filters are discussed on each individual report page description. See Web Reporting Page Descriptions.

Custom Reports

You can create a custom email security report pageby assembling charts (graphs) and tables from existing report pages.

Note On Email Security appliances, starting in release 9.6, “My Reports” is called “My Dashboard”.

To	Do This
Add modules to your custom report page	See: • Modules That Cannot Be Added to Custom Reports , on page 29 • Creating Your Custom Report Page , on page 29
View your custom report page	1 ChooseEmail > Reporting > My Reports. 2 Select the time range to viewThe time range selected applies to all reports, including all modules on the My Reports page. Newly-added modules appear at the top of the custom report.
Rearrange modules on your custom report page	Drag and drop modules into the desired location.
Delete modules from your custom report page	Click the [X] in the top right corner of the module.

To	Do This
Generate a CSV version of your custom report	See: • Generating Email Reports On Demand , on page 85
Periodically generate a CSV version of your custom report	See: • Scheduling Email Reports, on page 83

Modules That Cannot Be Added to Custom Reports

- All modules on the Management Appliance > Centralized Services > System Status page
- All modules on the Email > Reporting > Reporting Data Availability page
- All modules on the Email > Message Tracking > Message Tracking Data Availability page
- The following per-domain modules from the Sender Profile detail report page: Current Information from SenderBase, Sender Group Information, and Network Information
- The Past Year Virus Outbreak Summary chart and Past Year Virus Outbreaks table on the Outbreak Filters report page

Creating Your Custom Report Page

Before You Begin

- Ensure that the modules that you want to add can be added. See Modules That Cannot Be Added to Custom Reports , on page 29.
- Delete any default modules that you do not need by clicking the [X] in the top right corner of the module.

Step 1 Use one of the following methods to add a module to your custom report page:

Note Some modules are available only using one of these methods. If you cannot add a module using one method, try another method.

-
- Navigate to the report page under the Email tab that has the module that you want to add, then click the [+] button at the top of the module.
- Go toEmail > Reporting > My Reports, click the [+] Report Module button at the top of one of the sections, then select the report module that you want to add You may need to click the + button in each section on the My Reports page in order to find the module that you are looking for.

You can add each module only once; if you have already added a particular module to your report, the option to add it will not be available.

Step 2 If you add a module that you have customized (for example, by adding, deleting, or reordering columns , or by displaying non-default data in the chart), customize the modules on the My Reports page.
Modules are added with default settings. Time range of the original module is not maintained.

Step 3 If you add a chart that includes a separate legend (for example, a graph from the Overview page), add the legend separately. If necessary, drag and drop it into position beside the data it describes.

Viewing Details of Messages or Transactions Included in Reports

Step 1 Click any blue number in a table on a report page.
(Not all tables have these links.)

The messages or transactions included in that number are displayed in Message Tracking.

Step 2 Scroll down to see the list of messages or transactions.

What to Do Next

- Tracking Messages, on page 159

Improving Performance of Email Reports

If the performance of aggregated reporting decreases due to a large number of unique entries over the course of a month, use reporting filters to restrict the aggregation of data in reports that cover the previous year (Last Year reports). These filters can restrict detailed, individual IP, domain, or user data in reports. Overview reports and summary information remain available for all reports.

You can enable one or more of the reporting filters using the **reportingconfig > filters** menu in the CLI. The changes must be committed to take effect.

- **IP Connection Level Detail**. Enabling this filter prevents the Security Management appliance from recording information about individual IP addresses. This filter is appropriate for systems that process a large number of incoming IP addresses due to attacks.

  This filter affects the following Last Year reports:

  - Sender Profile for Incoming Mail
  - IP Addresses for Incoming Mail
  - IP Addresses for Outgoing Senders

- **User Detail**. Enabling this filter prevents the Security Management appliance from recording information about individual users sending and receiving mail and the content filters that are applied to the users'

mail. This filter is appropriate for appliances that process mail for millions of internal users or if the system does not validate recipient addresses.

This filter affects the following Last Year reports:

- Internal Users
- Internal User Details
- IP Addresses for Outgoing Senders
- Content Filters

- **Mail Traffic Detail**. Enabling this filter prevents the Security Management appliance from recording information about individual domains and networks that the appliances monitor. This filter is appropriate when the number of valid incoming or outgoing domains is measured in the tens of millions.

  This filter affects the following Last Year reports:

  - Domains for Incoming Mail
  - Sender Profile for Incoming Mail
  - Internal User Details
  - Domains for Outgoing Senders

Note To view up-to-the-minute reporting data for the preceding hour, you must log in to an individual appliance and view the data there.

Printing and Exporting Reporting Data

Table 7: Printing and Exporting Reporting Data

To Get This	CSV	Do This	Notes
Raw data See also Exporting Report Data as a Comma Separated Values (CSV) File , on page 32	•	Click the Export link below the chart or table.	The CSV file contains all applicable data, including the data visible in the chart or table.
	•	Create a scheduled or on-demand report. See: • Generating Email Reports On Demand , on page 85 • Scheduling Email Reports, on page 83	Each CSV file may contain up to 100 rows. If a report contains more than one table, a separate CSV file is created for each table. Some extended reports are not available in CSV format.

Exporting Report Data as a Comma Separated Values (CSV) File

You can export raw data to a comma-separated values (CSV) file, which you can access and manipulate using database applications such as Microsoft Excel. For different ways to export data, see Printing and Exporting Reporting Data , on page 31.

Because CSV exports include only raw data, exported data from a web-based report page may not include calculated data such as percentages, even if that data appears in the web-based report.

For email message tracking and reporting data, the exported CSV data will display all data in GMT regardless of what is set on the Security Management appliance. This simplifies using data independently from the appliance, particularly when referencing data from appliances in multiple time zones.

The following example is an entry from a raw data export of the Anti-Malware category report, where Pacific Daylight Time (PDT) is displayed as GMT - 7 hours:

Begin Timestamp, End Timestamp, Begin Date, End Date, Name, Transactions Monitored, Transactions Blocked, Transactions Detected

1159772400.0, 1159858799.0, 2006-10-02 07:00 GMT, 2006-10-03 06:59 GMT, Adware, 525, 2100, 2625

Table 8: Viewing Raw Data Entries

Category Header	Value	Description
Begin Timestamp	1159772400.0	Query start time in number of seconds from epoch.
End Timestamp	1159858799.0	Query end time in number of seconds from epoch.
Begin Date	2006-10-02 07:00 GMT	Date the query began.
End Date	2006-10-03 06:59 GMT	Date the query ended.
Name	Adware	Name of the malware category.
Transactions Monitored	525	Number of transactions monitored.
Transactions Blocked	2100	Number of transactions blocked.
Transactions Detected	2625	Total number of transactions: Number of transactions detected + Number of transactions blocked.

Note Category headers are different for each type of report.If you export localized CSV data, the headings may not be rendered properly in some browsers. This occurs because some browsers may not use the proper character set for the localized text. To work around this problem, you can save the file to your local machine, and open the file on any web browser using **File > Open**. When you open the file, select the character set to display the localized text.

Subdomains vs. Second Level Domains in Reporting and Tracking

In reporting and tracking searches, second-level domains (regional domains listed at http://george.surbl.org/two-level-tlds) are treated differently from subdomains, even though the two domain types may appear to be the same. For example:

- Reports will not include results for a two-level domain such as co.uk , but will include results for foo.co.uk . Reports include subdomains under the main corporate domain, such as cisco.com .
- Tracking search results for the regional domain co.uk will not include domains such as foo.co.uk , while search results for cisco.com will include subdomains such as subdomain.cisco.com .

Troubleshooting All Reports

Unable to View Report Data on Backup Security Management Appliance

Problem

You are unable to select a single Email Security appliance for which to view report data. The **View Data For** option does not appear on the reporting page.

Solution

See also Availability of Services During Backups , on page 292.

Reporting Is Disabled

Problem

Canceling a backup in progress can disable reporting.

Solution

Reporting functionality will be restored after a backup is completed.

Email and Web Reports

For information specific to Email reports, see Using Centralized Email Security Reporting on the Legacy Web Interface, on page 35.

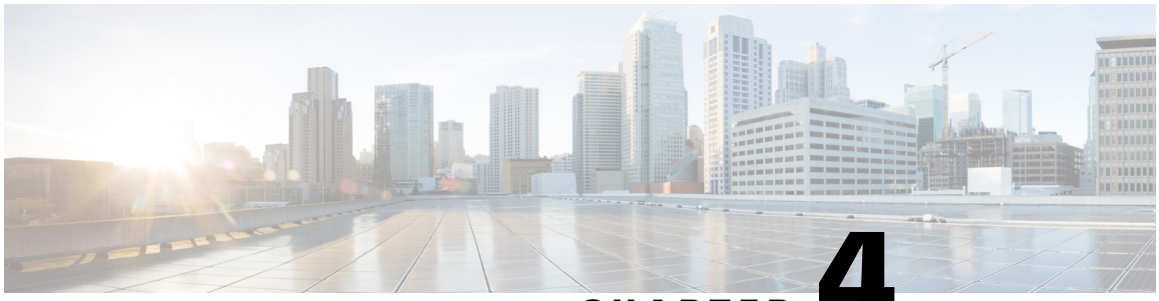

CHAPTER 4

Using Centralized Email Security Reporting on the Legacy Web Interface

This chapter contains the following sections:

Centralized Email Reporting Overview

Your Cisco Content Security Management appliance shows aggregated information from individual or multiple Email Security appliances so that you can monitor your email traffic patterns and security risks. You can run reports in real-time to view an interactive display of system activity over a specific period of time, or you can schedule reports and run them at regular intervals. Reporting functionality also allows you to export raw data to a file.

This feature centralizes the reports listed under the Monitor menu of the Email Security appliance.

The Centralized Email Reporting feature not only generates high-level reports, allowing you to understand what is happening on their network, but it also allows you to drill down and see traffic details for a particular domain, user, or category.

The Centralized Tracking feature allows you to track email messages that traverse multiple Email Security appliances.

Note The Email Security appliance only stores data if local reporting is used. If centralized reporting is enabled for the Email Security appliance then the Email Security appliance does NOT retain any reporting data except for System Capacity and System Status. If Centralized Email Reporting is not enabled, the only reports that are generated are System Status and System Capacity.

For more information about availability of report data during and after the transition to centralized reporting, see the "Centralized Reporting Mode" section of the documentation or online help for your Email Security appliance.

Setting Up Centralized Email Reporting

To set up centralized email reporting, complete the following procedures in order:

Note If reporting and tracking are not consistently and simultaneously enabled and functioning properly, or are not consistently and simultaneously either centralized or stored locally on each Email Security appliance, then the message tracking results when drilling down from reports will not match expected results. This is because the data for each feature (reporting, tracking) is captured only while that feature is enabled.

Enabling Centralized Email Reporting on the Security Management Appliance

Before You Begin

- All Email Security appliances should be configured and working as expected before you enable centralized reporting.
- Before enabling centralized email reporting, ensure that sufficient disk space is allocated to that service. See the Managing Disk Space , on page 333.

Step 1 On the Security Management appliance, choose **Management Appliance > Centralized Services > Email > Centralized Reporting**.

Step 2 Click **Enable**.

Step 3 If you are enabling centralized email reporting for the first time after running the System Setup Wizard, review the end user license agreement, and click **Accept**.

Step 4 Submit and commit your changes.

Note If you have enabled email reporting on the appliance, and there is no disk space allocated for this action, centralized email reporting will not work until disk space is allocated. As long as the quota you are setting the Email Reporting and Tracking to is larger than the currently used disk space, you will not lose any reporting and tracking data. See the Managing Disk Space , on page 333 section, for more information.

Adding the Centralized Email Reporting Service to Each Managed Email Security Appliance

The steps you follow depend on whether or not you have already added the appliance while configuring another centralized management feature.

Step 1 Choose **Management Appliance > Centralized Services > Security Appliances**.

Step 2 If you have already added the Email Security appliance to the list on this page:

a) Click the name of an Email Security appliance.

b) Select the **Centralized Reporting** service.

Step 3 If you have not yet added Email Security appliances:

a) Click Add Email Appliance.

b) In the Appliance Name and IP Address text fields, type the appliance name and the IP address for the Management interface of the Security Management appliance.

Note If you enter A DNS name in the IP Address text field, it will be immediately resolved to an IP address when you click **Submit**.

c) The Centralized Reporting service is pre-selected.

d) Click **Establish Connection.**

e) Enter the user name and password for an administrator account on the appliance to be managed, then click **Establish Connection.**.

Note You enter the login credentials to pass a public SSH key for file transfers from the Security Management appliance to the remote appliance. The login credentials are not stored on the Security Management appliance.

f) Wait for the Success message to appear above the table on the page.

g) Click **Test Connection.**.

h) Read test results above the table.

Step 4 Click **Submit**.

Step 5 Repeat this procedure for each Email Security appliance for which you want to enable Centralized Reporting.

Step 6 Commit your changes.

Creating Email Reporting Groups

You can create groups of Email Security appliances for which to view reporting data from the Security Management appliance.

A group can include one or more appliances, and an appliance may belong to more than one group.

Before You Begin

Make sure centralized reporting is enabled for each appliance. See Adding the Centralized Email Reporting Service to Each Managed Email Security Appliance , on page 103.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Centralized Reporting**.

Step 3 Click **Add Group** .

Step 4 Enter a unique name for the group.
The Email Security appliance list displays the Email Security appliances that you added to the Security Management appliance. Select the appliances that you want to add to the group.

The maximum number of groups that can be added is smaller than or equal to the maximum number of email appliances that can be connected.

Note If you added an Email Security appliance to the Security Management appliance, but you do not see it in the list, edit the configuration of the Email Security appliance so that the Security Management appliance is collecting reporting data from it.

Step 5 Click **Add** to add the appliances to the Group Members list.

Step 6 Submit and commit your changes.

Enabling Centralized Email Reporting on Email Security Appliances

You must enable centralized email reporting on each managed Email Security appliance appliance.

For instructions, see the “Configuring an Email Security Appliance to Use Centralized Reporting” section of the documentation or online help for your Email Security appliance.

Working with Email Report Data

- For options for accessing and viewing report data, see Ways to View Reporting Data , on page 23.
- To customize your view of report data, see Customizing Your View of Report Data , on page 25
- To search for specific information within your data, see Searching and the Interactive Email Report Pages , on page 39.
- To print or export report information, see Printing and Exporting Reporting Data , on page 31
- To understand the various interactive report pages, see Understanding the Email Reporting Pages, on page 40.
- To generate a report on demand, see Generating Email Reports On Demand , on page 85.
- To schedule reports to run automatically at intervals and times that you specify, see Scheduling Email Reports, on page 83.
- To view archived on-demand and scheduled reports, see Viewing and Managing Archived Email Reports , on page 87.

- For background information, How the Security Management Appliance Gathers Data for Reports, on page 24.
- To improve performance when working with large amounts of data, see Improving Performance of Email Reports , on page 30.
- To get details about an entity or number that appears as a blue link in a chart or table, click the entity or number.

  For example, if your permissions allow you to do so, you can use this feature to view details about messages that violate Content Filtering or Data Loss Prevention policies. This performs the relevant search in Message Tracking. Scroll down to view results.

Searching and the Interactive Email Report Pages

Many of the interactive email reporting pages include a '**Search For:**' drop-down menu at the bottom of the page.

From the drop-down menu, you can search for several types of criteria, including the following:

- IP address
- Domain
- Network owner
- Internal user
- Destination domain
- Internal sender domain
- Internal sender IP address
- Incoming TLS domain
- Outgoing TLS domain
- SHA-256

For most searches, choose whether to exactly match the search text or look for items starting with the entered text (for example, starts with "ex" will match "example.com").

For IPv4 searches, the entered text is always interpreted as the beginning of up to four IP octets in dotted decimal format. For example, '17.*' will search in the range 17.0.0.0 through 17.255.255.255, so it will match 17.0.0.1 but not 172.0.0.1. For an exact match search, enter all four octets. IP address searches also support Classless Inter-Domain Routing (CIDR) format (17.16.0.0/12).

For IPv6 searches, you can enter addresses using the formats in the following examples:

- 2001:db8:2004:4202::0-2001:db8:2004:4202::ff
- 2001:db8:2004:4202::
- 2001:db8:2004:4202::23
- 2001:db8:2004:4202::/64

Understanding the Email Reporting Pages

Note This list represents the reports available in the latest supported release of AsyncOS for Email Security appliances. If your Email Security appliances are running earlier releases of AsyncOS, not all of these reports are available.

Table 9: Email Reporting Tab Options

Email Reporting Menu	Action
Email Reporting Overview Page	The Overview page provides a synopsis of the activity on your Email Security appliances. It includes graphs and summary tables for the incoming and outgoing messages. For more information, see the Email Reporting Overview Page, on page 46.
Incoming Mail Page	The Incoming Mail page provides interactive reporting on the real-time information for all remote hosts connecting to your managed Email Security appliances. You can gather information about the IP addresses, domains, and network owners (organizations) sending mail to your system. For more information, see the Incoming Mail Page, on page 51.
Sender Groups Report Page	The Sender Groups report page provides a summary of connections by sender group and mail flow policy action, allowing you to review SMTP connection and mail flow policy trends. For more information, see the Sender Groups Report Page, on page 54.
Outgoing Destinations Page	The Outgoing Destinations page provides information about the domains that your organization sends mail to. The top of the page includes graphs depicting the top destinations by outgoing threat messages and top destinations by outgoing clean messages. The bottom of the page displays a chart with columns sorted by total recipients (default setting). For more information, see the Outgoing Destinations Page, on page 55.
Outgoing Senders Page	The Outgoing Senders page provides information about the quantity and type of mail being sent from IP addresses and domains in your network. For more information, see the Outgoing Senders Page, on page 56.
Internal Users Page	The Internal Users provides information about the mail sent and received by your internal users *per email address*. A single user can have multiple email addresses. The email addresses are not combined in the report. For more information, see the Internal Users Page, on page 57.

Email Reporting Menu	Action
DLP Incidents	The DLP Incident Summary page shows information on the incidents of data loss prevention (DLP) policy violations occurring in outgoing mail. For more information, see the DLP Incidents , on page 59.
Message Filters	The Message Filters page shows information about the top message filter matches (which message filters had the largest number of matching messages) for incoming and outgoing messages. For more information, see the Message Filters , on page 60
Geo Distribution	The Geo Distribution page shows the: • Top incoming mail connections based on country of origin in graphical format. • Total incoming mail connections based on country of origin in tabular format. For more information, see the Geo Distribution, on page 60.
High Volume Mail	The High Volume Mail page identifies attacks involving a large number of messages from a single sender, or with identical subjects, within a moving one-hour period. For more information, see the High Volume Mail , on page 61.
Content Filters Page	The Content Filters page shows information about the top incoming and outgoing content filter matches (which content filter had the most matching messages). This page also displays the data as both bar charts and listings. Using the Content Filters page, you can review your corporate policies on a per-content-filter or per-user basis. For more information, see the Content Filters Page, on page 61.
DMARC Verification	The DMARC Verification page shows the top sender domains that failed Domain-based Message Authentication, Reporting and Conformance (DMARC) verification, and a summary of the actions taken for incoming messages from each domain. For more information, see the DMARC Verification , on page 62.
Macro Detection	The Macro Detection Report page shows the top incoming and outgoing macro-enabled attachments by file type detected by the content or message filters. For more information, see the Macro Detection, on page 62

Email Reporting Menu	Action
Virus Types Page	The Virus Types page provides an overview of the viruses that are sent to and from your network. The Virus Types page displays the viruses that have been detected by the virus scanning engines running on the Email Security appliances and are displayed on the Security Management appliance. Use this report to take action against a particular virus. For more information, see the Virus Types Page, on page 62.
URL Filtering Page	Use this page to view the URL categories most frequently occurring in messages, the most common URLs in spam messages, and the number of malicious and neutral URLs seen in messages. For more information, see the URL Filtering Page , on page 63.
Web Interaction Tracking Page	Identifies the end users who clicked URLs rewritten by policy or Outbreak Filter, and the action associated with each user click. For more information, see the Web Interaction Tracking Page, on page 64.
Forged Email Detection Page	The Forged Email Detection page includes the following reports: • **Top Forged Email Detection**. Displays the top ten users in the content dictionary that matched the forged From: header in the incoming messages. • **Forged Email Detection: Details**. Displays a list of all the users in the content dictionary that matched the forged From: header in the incoming messages and for a given user, the number of messages matched. See Forged Email Detection Page, on page 65.
Advanced Malware Protection (File Reputation and File Analysis) Reporting Pages	There are three reporting pages showing file reputation and analysis data. For more information, see the Advanced Malware Protection (File Reputation and File Analysis) Reporting Pages , on page 65.
Mailbox Auto Remediation	Use this page to view the details of the mailbox remediation results. See Mailbox Auto Remediation, on page 70
TLS Connections Page	The TLS Connections page shows the overall usage of TLS connections for sent and received mail. The report also shows details for each domain sending mail using TLS connections. For more information, see the TLS Connections Page, on page 71.
Inbound SMTP Authentication Page	The Inbound SMTP authentication page shows the use of client certificates and the SMTP AUTH command to authenticate SMTP sessions between the Email Security appliance and users' mail clients. For more information, see Inbound SMTP Authentication Page, on page 72.

Email Reporting Menu	Action
Outbreak Filters Page	The Outbreak Filters page shows information about recent outbreaks and the messages quarantined by Outbreak Filters. Use this page to monitor your defense against virus attacks. For more information, see the Outbreak Filters Page, on page 73.
Rate Limits Page	The Rate Limits page shows the mail senders (based on MAIL-FROM address) who exceed the threshold you set for the number of message recipients per sender. For more information, see the Rate Limits Page , on page 73.
System Capacity Page	Allows you to view the overall workload that is sending reporting data to the Security Management appliance. For more information, see the System Capacity Page, on page 76.
Reporting Data Availability Page	Allows you to get a glimpse of the impact of the reporting data on the Security Management appliance for each appliance. For more information, see the Reporting Data Availability Page, on page 79.
Scheduling Email Reports	Allows you to schedule reports for a specified time range. For more information, see the Scheduling Email Reports, on page 83.
Viewing and Managing Archived Email Reports	Allows you to view and manage archived reports. For more information, see the Viewing and Managing Archived Email Reports , on page 87. Also allows you to generate on-demand reports. See Generating Email Reports On Demand , on page 85.

Table Column Descriptions for Email Reporting Pages

Table 10: Table Column Descriptions for Email Reporting Pages

Column Name	
Incoming Mail Details	
Connections Rejected	All connections blocked by HAT policies. When the appliance is under heavy load, an exact count of rejected connections is not maintained on a per-sender basis. Instead, rejected connections counts are maintained only for the most significant senders in each time interval.
Connections Accepted	All connections accepted,
Total Attempted	All accepted and blocked connections attempted.

Column Name	
Stopped by Recipient Throttling	This is a component of Stopped by Reputation Filtering. It represents the number of recipient messages stopped because any of the following HAT limits have been exceeded: maximum recipients per hour, maximum recipients per message, or maximum messages per connection. This is summed with an estimate of the recipient messages associated with rejected or TCP refused connections to yield Stopped by Reputation Filtering.
Stopped by Reputation Filtering	The value for Stopped by Reputation Filtering is calculated based on several factors: • Number of "throttled" messages from this sender • Number of rejected or TCP refused connections (may be a partial count) • A conservative multiplier for the number of messages per connection When the appliance is under heavy load, an exact count of rejected connections is not maintained on a per-sender basis. Instead, rejected connections counts are maintained only for the most significant senders in each time interval. In this situation, the value shown can be interpreted as a "floor"; that is, at least this many messages were stopped. **Note** The Stopped by Reputation Filtering total on the Overview page is always based on a complete count of all rejected connections. Only the per-sender connection counts are limited due to load.
Stopped as Invalid Recipients	All mail recipients rejected by conversational LDAP rejection plus all RAT rejections.
Spam Detected	Any spam that has been detected.
Virus Detected	Any viruses that have been detected
Stopped by Content Filter	The total count of messages that were stopped by a content filter.
Total Threat	Total number of threat messages (stopped by reputation, stopped as invalid recipient, spam, plus virus)
Marketing	Number of messages detected as unwanted marketing messages.

Column Name	
Clean	All clean messages. Messages processed on appliances on which the graymail feature is not enabled are counted as clean.
User Mail Flow Details (Internal Users Page)	
Incoming Spam Detected	All incoming spam that is detected
Incoming Virus Detected	The incoming virus that has been detected.
Incoming Content Filter Matches	The incoming content filter matches that have been detected.
Incoming Stopped by Content Filter	The Incoming messages that were stopped due to content filters that have been set.
Incoming Clean	All incoming clean messages.
Outgoing Spam Detected	The outgoing spam that was detected.
Outgoing Virus Detected	The outgoing viruses that have been detected.
Outgoing Content Filter Matches	The outgoing content filter matches that have been detected.
Outgoing Stopped by Content Filter	The outgoing messages that were stopped due to content filters that have been set.
Outgoing Clean	All outgoing clean messages.
Incoming and Outgoing TLS Connections: TLS Connections Page	
Required TLS: Failed	All required TLS connections that failed.
Required TLS: Successful	All required TLS connections that are successful.
Preferred TLS: Failed	All preferred TLS connections that failed.
Preferred TLS: Successful	All preferred TLS connections that are successful.
Total Connections	Total number of TLS connections.
Total Messages	The total number of TLS messages.
Outbreak Filters	
Outbreak Name	The name of the outbreak.

Column Name	
Outbreak ID	The outbreak ID.
First Seen Globally	The first time the virus has been seen globally.
Protection Time	The time the virus has been protected.
Quarantined Messages	Messages related to the quarantine.

Email Reporting Overview Page

The **Email > Reporting > Overview** page on the Security Management appliance provides a synopsis of the email message activity from your Email Security appliances. The Overview page includes graphs and summary tables for the incoming and outgoing messages.

At a high level the **Overview** page shows you the incoming and outgoing mail graphs, and well as incoming and outgoing mail summaries.

The mail trend graphs provide a visual representation of the mail flow. You can use the mail trend graphs on this page to monitor the flow of all mail into and out of your appliances.

Note The Domain-Based Executive Summary Report and the Executive Summary report are based on the Email Reporting Overview Page, on page 46. For more information, see the Domain-Based Executive Summary Report, on page 81 and Executive Summary Report , on page 83

Table 11: Details on the Email Reporting Overview Page

Section	Description
Time Range	A drop-down list with options for choosing a time range to view. For more information, see the Choosing a Time Range for Reports , on page 26.
View Data for	Choose an Email Security appliance for which you want to view Overview data, or choose All Email Appliances. See also Viewing Reporting Data for an Appliance or Reporting Group, on page 26.

How Incoming Mail Messages are Counted

Counts of incoming messages are dependent on the number of recipients per message. For example, an incoming message from example.com sent to three recipients is counted as three messages coming from that sender.

Because the messages blocked by sender reputation filtering do not actually enter the work queue, the appliance does not have access to the list of recipients for an incoming message. In this case, a multiplier is used to estimate the number of recipients. This multiplier is based on research of a large sampling of existing customer data.

How Email Messages Are Categorized by the Appliances

As messages proceed through the email pipeline, they can apply to multiple categories. For example, a message can be marked as spam or virus positive; it can also match a content filter. The precedence of the various filters and scanning activities greatly impacts the results of message processing.

In the example above, the various verdicts follow these rules of precedence:

- Spam positive
- Virus positive
- Matching a content filter

Following these rules, if a message is marked as spam positive, and your anti-spam settings are set to drop spam positive messages, the message is dropped and the spam counter is incremented.

Further, if your anti-spam settings are set to let the spam positive message continue on in the email pipeline, and a subsequent content filter drops, bounces, or quarantines the message, the spam count is still incremented. The content filter count is only incremented if the message is not spam or virus positive.

Alternately, if the message were quarantined by Outbreak Filters, it would not be counted until it was released from the quarantine and again processed through the work queue.

For complete information about message processing precedence, see chapter about the email pipeline in the online help or user guide for your Email Security appliance.

Categorizing Email Messages on the Overview Page

Messages reported in the Incoming Mail Summary on the Overview report page are categorized as follows:

Table 12: Email Categories on Overview Page

Category	Description
Stopped by Reputation Filtering	All connections blocked by HAT policies multiplied by a fixed multiplier (see the How Incoming Mail Messages are Counted, on page 46) plus all recipients blocked by recipient throttling. The value for Stopped by Reputation Filtering is calculated based on several factors: • Number of “throttled” messages from this sender • Number of rejected or TCP refused connections (may be a partial count) • A conservative multiplier for the number of messages per connection When the appliance is under heavy load, an exact count of rejected connections is not maintained on a per-sender basis. Instead, rejected connections counts are maintained only for the most significant senders in each time interval. In this situation, the value shown can be interpreted as a “floor”; that is, at least this many messages were stopped. The Stopped by Reputation Filtering total on the Overview page is always based on a complete count of all rejected connections. Only the per-sender connection counts are limited due to load.
Invalid Recipients	All mail recipients rejected by conversational LDAP rejection plus all RAT rejections.
Spam Messages Detected	The total count of messages detected by the anti-spam scanning engine as positive or suspect. Additionally, messages that are both spam and virus positive.

Category	Description
Virus Messages Detected	The total count and percentage of messages detected as virus positive and not also spam. The following messages are counted in the "Virus Detected" category: • Messages with a virus scan result of "Repaired" or "Infectious" • Messages with a virus scan result of "Encrypted" when the option to count encrypted messages as containing viruses is selected • Messages with a virus scan result of "Unscannable" when the action for unscannable messages is NOT "Deliver" • Messages with a virus scan result of "Unscannable" or "Encrypted" when the option to deliver to an alternate mail host or an alternate recipient is selected • Messages that are deleted from the Outbreak quarantine, either manually or by timing out.
Detected by Advanced Malware Protection	A message attachment was found to be malicious by file reputation filtering. This value does not include verdict updates or files found to be malicious by file analysis.
Messages with Malicious URLs	One or more URLs in the message were found to be malicious by URL filtering.
Stopped by Content Filter	The total count of messages that were stopped by a content filter. If your access privileges allow you to view Message Tracking data: To view Message Tracking details for the Content Filter violations in this report, click a blue number link in the table.
Stopped by DMARC	The total count of messages that failed DMARC verification.
S/MIME Verification/Decryption Failed	The total count of messages that failed S/MIME verification, decryption, or both.
Marketing Messages	The total count of advertising messages sent by recognized professional marketing groups, for example Amazon.com. This list item appears on the page only if marketing data are present in the system. This number includes marketing messages identified both by Email Security appliances on which the graymail feature is enabled and by appliances on which Marketing Email Scanning under anti-spam settings is enabled.

Category	Description
Social Networking Messages	The total count of notification messages from social networks, dating websites, forums, and so on. Examples include LinkedIn and CNET forums. This information is determined by the graymail feature.
Bulk Messages	The total count of advertising messages sent by unrecognized marketing groups, for example, TechTarget, a technology media company. This information is determined by the graymail feature.
Graymail Messages	This number includes marketing messages detected by the graymail feature, plus social networking messages and bulk mail. It does not include marketing messages identified on appliances on which the graymail feature is not enabled, even if those totals are included in the Marketing Messages value. Click on the number corresponding to any of the graymail categories to view a list of messages belonging to that category using Message Tracking. See also Reporting of Graymail , on page 75.
S/MIME Verification/Decryption Successful	The total count of messages that were successfully verified, decrypted, or decrypted and verified using S/MIME.
Clean Messages Accepted	This category is mail that is accepted and deemed to be virus and spam free. The most accurate representation of clean messages accepted when taking per-recipient scanning actions (such as splintered messages being processed by separate mail policies) into account. However, because messages that are marked as spam or virus positive and still delivered are not counted, the actual number of messages delivered may differ from the clean message count. If messages match a *message filter* and are not dropped or bounced by the filter, they are treated as clean. Messages dropped or bounced by a message filter are not counted in the totals. Messages processed on appliances on which the graymail feature is not enabled are counted as clean.
Total Attempted Messages	This number includes spam, marketing messages (whether found by the graymail feature or by Marketing Email Scanning functionality in the anti-spam feature), social networking messages, bulk mail, and clean messages.

Note If you have configured your anti-virus settings to deliver unscannable or encrypted messages, these messages will be counted as clean messages and not virus positive. Otherwise, the messages are counted as virus positive.Additionally, if messages match a *message filter* and are not dropped or bounced by the filter, they are treated as clean. Messages dropped or bounced by a message filter are not counted in the totals.

Incoming Mail Page

The **Email > Reporting > Incoming Mail**page on the Security Management appliance provides interactive reporting on the real-time information for all remote hosts connecting to your managed Security Management appliances. You can gather information about the IP addresses, domains, and network owners (organizations) sending mail to your system. You can also perform a Sender Profile search on IP addresses, domains, or organizations that have sent mail to you.

The Incoming Mail Details interactive table displays detailed information about the particular IP address, domain, or network owner (organization). You can access a Sender Profile page for any IP address, domain, or network owner by clicking the corresponding link at the top of the **Incoming Mail** page, or on other Sender Profile pages.

From the Incoming Mail pages you can:

- Perform a search on IP addresses, domains, or network owners (organizations) that have sent mail to your Security Management appliances. See Searching and the Interactive Email Report Pages , on page 39.
- View the Sender Groups report to monitor connections according to the specific sender group and mail flow policy actions. See the Sender Groups Report Page, on page 54 for more information.
- See detailed statistics on senders that have sent mail to your appliances. The statistics include the number of attempted messages broken down by security service (sender reputation filtering, anti-spam, anti-virus, and so forth).
- Sort by senders who have sent you a high volume of spam or virus email, as determined by anti-spam or anti-virus security services.
- Use the SenderBase Reputation Service to examine the relationship between specific IP addresses, domains, and organizations to obtain information about a sender.
- Obtain more information about a sender from the SenderBase Reputation Service, including a sender's SenderBase Reputation Score (SBRS) and which sender group the domain matched most recently. Add senders to sender groups.
- Obtain more information about a specific sender who has sent a high volume of spam or virus email, as determined by the anti-spam or anti-virus security services.

Views Within the Incoming Mail Page

The **Incoming Mail** page has three different views:

- IP Addresses
- Domains

- Network Owners

These views provide a snapshot of the remote hosts connecting to the system in the context of the selected view.

Additionally, in the Incoming Mail Details section of the Incoming Mail Page, you can click on a Sender's IP Address, Domain name, or Network Owner Information to retrieve specific Sender Profile Information. For more information on Sender Profile information, see the Sender Profile Pages, on page 53.

Note *Network owners* are entities that contain domains. *Domains* are entities that contain IP addresses.

Depending on the view you select, the Incoming Mail Details interactive table displays the top IP addresses, domains, or network owners that have sent mail to all public listeners configured on the Email Security appliances. You can monitor the flow of all mail into your appliances.

Click an IP address, domain, or network owner to access details about the sender on the Sender Profile page. The Sender Profile page is an Incoming Mail page that is specific to a particular IP address, domain, or network owner.

To access the mail flow information by sender group, click the **Sender Groups Report** link at the bottom of the Incoming Mail page. See Sender Profile Pages, on page 53.

In some cases, some of the report pages contain several unique sub-reports that can be accessed from the top-level page. For example, the Incoming Mail report page on the Security Management appliance allows you to see information for individual IP Addresses, Domains and Network Owners. Each of these are sub-pages are accessed from the Incoming Mail report page.

Results for each of these sub-report pages are generated on one consolidated report when you click on the Printable PDF link at the top-right of the top-level page; in this case the Incoming Mail report page. See important information in Understanding the Email Reporting Pages, on page 40.

The **Email > Reporting > Incoming Mail** page offers the following views: **IP Addresses**, **Domains**, or **Network Owners**

See the Incoming Mail Details Table, on page 53 for an explanation of the data included in the Incoming Mail Details interactive table.

From the **Incoming Mail** page you can also generate a PDF or export raw data to a CSV file. For information on printing or exporting a file, see the Understanding the Email Reporting Pages, on page 40.

Note You can generate a scheduled report for the Incoming Mail report page. See the Scheduling Email Reports, on page 83.

"No Domain Information" Link

Domains that have connected to the Security Management appliances and could not be verified with a double-DNS lookup are automatically grouped into the special domain called "No Domain Information." You can control how these types of unverified hosts are managed via Sender Verification. For more information about Sender Verification, see the documentation or online help for your Email Security appliance.

You can use the Items Displayed menu to select the number of senders to display in the list.

Time Ranges in the Mail Trend Graphs

You can select varying degrees of granularity to see your data in a mail graph. You can select a day, week, month, and year views of the same data. Because the data is monitored in real time, information is periodically updated and summarized in the database.

For more information on time ranges, see Choosing a Time Range for Reports , on page 26.

Incoming Mail Details Table

The interactive Incoming Mail Details table at the bottom of the **Incoming Mail** page lists the top senders that have connected to public listeners on the Email Security appliances. The table shows domains, IP addresses, or network owners, based on the view selected. Click the column headings to sort the data.

The system acquires and verifies the validity of the remote host's IP address by performing a *double DNS lookup* . For more information about double DNS lookups and sender verification, see the documentation or online help for your Email Security appliance.

For senders, that is Network Owner, IP Address or Domain, listed in the first column of the Incoming Mail Details table, or on the Top Senders by Total Threat Messages, click the **Sender** or **No Domain Information** link to view more information about the sender. The results appear on a **Sender Profile** page, which includes real-time information from the SenderBase Reputation Service. From the Sender Profile page, you can view for more information about specific IP addresses or network owners. For more information, see the Sender Profile Pages, on page 53.

You can also view the Sender Groups report, by clicking **Sender Groups report** at the bottom of the Incoming Mail page. For more information about the Sender Groups report page, see the Sender Groups Report Page, on page 54.

If your access privileges allow you to view Message Tracking data: To view Message Tracking details for the Content Filter violations in this report, click a blue number link in the table.

Sender Profile Pages

When you click a sender in the Incoming Mail Details interactive table, on the **Incoming Mail** page, the Sender Profile page appears. It shows detailed information about the particular IP address, domain, or network owner (organization). You can access a Sender Profile page for any IP address, domain, or network owner by clicking the corresponding link on the Incoming Mail page or on other Sender Profile pages.

Network owners are entities that contain domains. *Domains* are entities that contain IP addresses.

The Sender Profile pages displayed for IP addresses, domains, and network owners vary slightly. For each, the page contains a graph and summary table for incoming mail from the particular sender. Below the graph, a table lists the domains or IP addresses associated with the sender. (The Sender Profile page for an individual IP address does not contain a more granular listing.) The Sender Profile page also includes an information section with the current SenderBase, sender group, and network information for the sender.

- Network Owner profile pages contain information for the network owner, as well as the domains and IP addresses associated with that network owner.
- Domain profile pages contain information for the domains and IP addresses associated with that domain.
- IP address profile pages contain information about the IP address only.

Each Sender Profile page contains the following data in the Current Information table at the bottom of the page:

- The global information from the SenderBase Reputation Service, including:
  - IP address, domain name, and/or network owner
  - Network owner category (network owner only)
  - CIDR range (IP addresses only)
  - Daily magnitude and monthly magnitude for the IP address, domain, and/or network owner
  - Days since the first message was received from this sender
  - Last sender group and whether DNS verified (IP address sender profile page only)

  Daily magnitude is a measure of how many messages a domain has sent over the last 24 hours. Similar to the Richter scale used to measure earthquakes, SenderBase magnitude is a measure of message volume calculated using a log scale with a base of 10. The maximum theoretical value of the scale is set to 10, which equates to 100% of the world's email message volume. Using the log scale, a one-point increase in magnitude equates to a 10x increase in actual volume.

  Monthly magnitude is calculated using the same approach as daily magnitude, except the percentages are calculated based on the volume of email sent over the last 30 days.

  - Average magnitude (IP addresses only)
  - Lifetime volume / 30 day volume (IP address profile pages only)
  - Bonded sender status (IP address profile pages only)
  - SenderBase Reputation Score (IP address profile pages only)
  - Days since first message (network owner and domain profile pages only)
  - Number of domains associated with this network owner (network owner and domain profile pages only)
  - Number of IP addresses in this network owner (network owner and domain profile pages only)
  - Number of IP addresses used to send email (network owner pages only)

  Click **More from SenderBase** to see a page with all information supplied by the SenderBase Reputation Service.

- Details about the domains and IP addresses controlled by this network owner appear on network owner profile pages. Details about the IP addresses in the domain appear on domain pages.

From a domain profile page, you can click on a specific IP address to view specific information, or view an organization profile page.

Sender Groups Report Page

The **Sender Groups report** page provides a summary of connections by sender group and mail flow policy action, allowing you to review SMTP connection and mail flow policy trends. The Mail Flow by Sender Group listing shows the percentage and number of connections for each sender group. The Connections by Mail Flow Policy Action chart shows the percentage of connections for each mail flow policy action. This page provides an overview of the effectiveness of your Host Access Table (HAT) policies. For more information about the HAT, see the documentation or online help for your Email Security appliance.

To view the Sender Groups report page, select **Email > Reporting > Sender Groups**.

From the **Sender Group Report** page you can also generate a PDF or export raw data to a CSV file. For information on printing or exporting a file, see the Understanding the Email Reporting Pages, on page 40.

Note You can generate a scheduled report for the Sender Group report page. See the Scheduling Email Reports, on page 83.

Outgoing Destinations Page

The **Email > Reporting > Outgoing Destinations** page provides information about the domains that your organization sends mail to.

Use the Outgoing Destinations page to answer the following types of questions:

- Which domains are the Email Security appliances sending mail to?
- How much mail is sent to each domain?
- How much of that mail is clean, spam-positive, virus-positive, malware or stopped by a content filter?
- How many messages are delivered and how many messages are hard-bounced by the destination servers?

The following list explains the various sections on the **Outgoing Destinations** page:

Table 13: Details on the Email Reporting Outgoing Destinations Page

Section	Description
Time Range (drop-down list)	A drop-down list that can range from a day to 90 days or a custom range. For more information on time ranges and customizing this for your needs, see the Choosing a Time Range for Reports , on page 26.
Top Destination by Total Threat	The top destination domains of outgoing threat messages (spam, antivirus, etc.) sent by your organization. Total threat include threats that are spam or virus positive or that triggered a content filter.
Top Destination by Clean Messages	The top destination domains of clean outgoing messages sent by your organization.
Outgoing Destination Details	All details related to the destination domains of all outgoing messages sent by your organization, sorted by total recipients. Details include detected spam, viruses, clean messages etc. If your access privileges allow you to view Message Tracking data: To view Message Tracking details for the Content Filter violations in this report, click a blue number link in the table.

From the **Outgoing Destinations** page you can also generate a PDF or export raw data to a CSV file. For information on printing or exporting a file, see the Understanding the Email Reporting Pages, on page 40.

Note You can generate a scheduled report for the Outgoing Destinations page. See the Scheduling Email Reports, on page 83.

Outgoing Senders Page

The **Email > Reporting > Outgoing Senders** page provides information about the quantity and type of mail being sent from IP addresses and domains in your network.

Use the Outgoing Senders page to answer the following types of questions:

- Which IP addresses are sending the most virus-positive, or spam-positive or malware email?
- Which IP addresses trigger content filters the most frequently?
- Which domains are sending the most mail?
- What are the total number of recipients that are being processed where a delivery was attempted.

To view the **Outgoing Senders** page, perform the following:

You can see the results of the Outgoing senders with two types of views:

- **Domain**: This view allows you to see the volume of mail that is being sent by each domain
- **IP address**: This view allows you to see which IP addresses are sending the most virus messages or triggering content filters.

The following list explains the various sections on the **Outgoing Senders** page for both views:

Table 14: Details on the Email Reporting Outgoing Sender Page

Section	Description
Time Range (drop-down list)	A drop-down list that can range from a day to 90 days or a custom range. For more information on time ranges and customizing this for your needs, see the Choosing a Time Range for Reports , on page 26.
Top Senders by Total Threat Messages	The top senders (by IP address or domain) of outgoing threat messages (spam, antivirus, etc.) in your organization.
Top Sender by Clean Messages	The top senders (by IP address or domain) of clean outgoing messages sent in your organization.

Section	Description
Sender Details	All details on the senders (by IP address or domain) of all outgoing messages sent by your organization. Details include detected spam, viruses, clean messages, etc. If your access privileges allow you to view Message Tracking data: To view Message Tracking details for the DLP and Content Filter violations in this report, click a blue number link in the table.

Note This page does not display information about message delivery. To track delivery information, such as the number of messages from a particular domain that were bounced, log in to the appropriate Email Security appliance and choose **Monitor > Delivery Status**.

From the **Outgoing Senders** page you can also generate a PDF or export raw data to a CSV file. For information on printing or exporting a file, see the Understanding the Email Reporting Pages, on page 40.

Note You can generate a scheduled report for the **Outgoing Senders** page. See the Scheduling Email Reports, on page 83.

Internal Users Page

The **Email > Reporting > Internal Users** page provides information about the mail sent and received by your internal users *per email address*. A single user can have multiple email addresses. The email addresses are not combined in the report.

Use the Internal Users interactive report page to answer these types of questions:

- Who is sending the most external email?
- Who receives the most clean email?
- Who receives the largest number of graymail messages?
- Who receives the most spam?
- Who is triggering which content filters?
- Whose email is getting caught by content filters?

Table 15: Details on the Email Reporting Internal Users Page

Section	Description
Time Range (drop-down list)	A drop-down list that can range from a day to 90 days or a custom range. For more information on time ranges and customizing this for your needs, see the Choosing a Time Range for Reports , on page 26.

Section	Description
Top Users by Clean Incoming Messages	The top users by (by IP address or domain) of clean incoming messages sent in your organization.
Top Users by Clean Outgoing Messages	The top users (by IP address or domain) of clean outgoing messages sent in your organization.
User Mail Flow Details	The User Mail Flow Details interactive section breaks down the mail received and sent by each email address. You can sort the listing by clicking the column headers. To view details for a user, click the user name in the Internal User column. For more information, see the Internal User Details Page, on page 58. If your access privileges allow you to view Message Tracking data: To view Message Tracking details for the Content Filter violations in this report, click a blue number link in the table.

From the **Internal Users** page you can also generate a PDF or export raw data to a CSV file. For information on printing or exporting a file, see the Understanding the Email Reporting Pages, on page 40.

Note You can generate a scheduled report for the Internal Users page. See the Scheduling Email Reports, on page 83.

Internal User Details Page

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

Inbound Internal Users are the users for which you received email, based on the Rcpt To: address. Outbound Internal Users are based on the Mail From: address and are useful when tracking the types of email that senders on your internal network are sending.

Click a content filter name to view detailed information for that filter on the corresponding content filter information page (see Content Filters Page, on page 61). You can use this method to view a list of all users who sent or received mail that matched the particular content filter.

Note Some outbound mail (such as bounces) has a null sender. They are counted as outbound “unknown.”

Searching for a Specific Internal User

With the search form at the bottom of the User Mail Summary page and the User Mail Flow Details page, you can search for a specific internal user (email address). Select whether to exactly match the search text or look for items starting with the entered text (for example, starts with “ex” will match “example@example.com”).

DLP Incidents

The **Email > Reporting > DLP Incidents (DLP Incident Summary)** page shows information on the incidents of data loss prevention (DLP) policy violations occurring in outgoing mail. The Email Security appliance uses the DLP email policies enabled in the Outgoing Mail Policies table to detect sensitive data sent by your users. Every occurrence of an outgoing message violating a DLP policy is reported as an incident.

Using the DLP Incident Summary report, you can answer these kinds of questions:

- What type of sensitive data is being sent by your users?
- How severe are these DLP incidents?
- How many of these messages are being delivered?
- How many of these messages are being dropped?
- Who is sending these messages?

The DLP Incident Summary page contains two main sections:

- the DLP incident trend graphs summarizing the top DLP incidents by severity (Low, Medium, High, Critical) and policy matches,
- the DLP Incident Details listing

Table 16: Details on the Email Reporting DLP Incident Summary Page

Section	Description
Time Range (drop-down list)	A drop-down list that can range from a day to 90 days or a custom range. For more information on time ranges and customizing this for your needs, see the Choosing a Time Range for Reports , on page 26.
Top Incidents by Severity	The top DLP incidents listed by severity.
Incident Summary	The DLP policies currently enabled for each email appliance's outgoing mail policies are listed in the DLP Incident Details interactive table at the bottom of the **DLP Incident Summary** page. Click the name of a DLP policy to view more detailed information.
Top DLP Policy Matches	The top DLP Policies that have been matched.
DLP Incident Details	The DLP Incident Details table shows the total number of DLP incidents per policy, with a breakdown by severity level, and whether any of the messages were delivered in the clear, delivered encrypted, or dropped. For more information on the DLP Incidents Details table, see the DLP Incidents Details Table, on page 60.

Click the name of a DLP policy to view detailed information on the DLP incidents detected by the policy. You can use this method to get a list of users who sent mail that contained sensitive data detected by the policy.

DLP Incidents Details Table

The DLP Incident Details table is an interactive table that shows the total number of DLP incidents per policy, with a breakdown by severity level, and whether any of the messages were delivered in the clear, delivered encrypted, or dropped. Click the column headings to sort the data.

To find out more information about any of the DLP Policies listed in this table, click the name of the DLP Policy and the DLP Policy Page appears.For more information, see DLP Policy Detail Page, on page 60.

If your access privileges allow you to view Message Tracking data: To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

DLP Policy Detail Page

If you click on a name of a DLP policy in the DLP Incident Details table, the resulting DLP Policy Detail page displays the DLP incidents data for the policy. The page displays graphs on the DLP Incidents based by Severity.

The page also includes an Incidents by Sender table at the bottom of the page that lists each internal user who has sent a message that violated the DLP policy. The table also shows the total number of DLP incidents for this policy per user, with a breakdown by severity level, and whether any of the messages were delivered in the clear, delivered encrypted, or dropped. You can use the Incidents by Sender table to find out which users may be sending your organization's sensitive data to people outside your network.

Clicking the sender name on the incident detail page opens up the Internal Users page. See the Internal Users Page, on page 57 for more information.

Message Filters

The Message Filters page shows information about the top message filter matches (which message filters had the largest number of matching messages) for incoming and outgoing messages.

Geo Distribution

You can use the Geo Distribution report page to view:

- Top incoming mail connections based on country of origin in graphical format.
- Total incoming mail connections based on country of origin in tabular format.

The following are the scenarios when no country information is displayed for the top and total incoming mail connections:

- The sender IP address belongs to a private IP address
- The sender IP address does not get a valid SBRS.

High Volume Mail

Use reports on this page to:

- Identify attacks involving a large number of messages from a single sender, or with identical subjects, within a moving one-hour period.
- Monitor top domains to ensure that such attacks do not originate in your own domain. If this situation occurs, one or more accounts in your organization may be compromised.
- Help identify false positives so you can adjust your filters accordingly.

Reports on this page show data only from message filters that use the Header Repeats rule and that pass the number-of-messages threshold that you set in that rule. When combined with other rules, the Header Repeats rule is evaluated last, and is not evaluated at all if the message disposition is determined by a preceding condition. Similarly, messages caught by Rate Limiting never reach Header Repeats message filters. Therefore, some messages that might otherwise be considered high-volume mail may not be included in these reports. If you have configured your filters to whitelist certain messages, those messages are also excluded from these reports.

For more information about message filters and the Header Repeats rule, see the online help or user guide for your Email Security appliance.

Related Topics

- Rate Limits Page , on page 73

Content Filters Page

The **Email > Reporting > Content Filters** page shows information about the top incoming and outgoing content filter matches (which content filter had the most matching messages). The page displays the data as both bar charts and listings. Using the Content Filters page, you can review your corporate policies on a per-content-filter or per-user basis and answer the following types of questions:

- Which content filter is triggered the most by incoming or outgoing mail?
- Who are the top users sending or receiving mail that triggers a particular content filter?

To view more information about a specific filter, click the name of the filter. The Content Filter Details page appears. For more information on Content Filter details page, see the Content Filter Details Page, on page 62.

If your access privileges allow you to view Message Tracking data: To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

From the **Content Filters** page you can also generate a PDF or export raw data to a CSV file. For information on printing or exporting a file, see the Understanding the Email Reporting Pages, on page 40.

Note You can generate a scheduled report for the Content Filter page. See the Scheduling Email Reports, on page 83.

Content Filter Details Page

The Content Filter Detail page displays matches for the filter over time, as well as matches by internal user.

In the Matches by Internal User section, click the name of a user to view the detail page for the internal user (email address). For more information, see Internal User Details Page, on page 58.

If your access privileges allow you to view Message Tracking data: To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

DMARC Verification

The DMARC Verification page shows the top sender domains that failed Domain-based Message Authentication, Reporting and Conformance (DMARC) verification, and a summary of the actions taken for incoming messages from each domain. You can use this report to fine-tune your DMARC settings and answer these kinds of questions:

- Which domains sent the most messages that failed DMARC verification?
- For each domain, what actions were taken on messages that failed DMARC verification?

For more information about DMARC verification, see the Email Authentication chapter in the online help or user guide for your Email Security appliance.

Macro Detection

You can use the Macro Detection report page to view:

- Top Incoming Macro-Enabled Attachments by File Type in graphical and tabular format.
- Top Outgoing Macro-Enabled Attachments by File Type in graphical and tabular format.

You can click on the number of macro-enabled attachments to view the related messages in Message Tracking.

Note During report generation:

- If one or more macros are detected within an archive file, the Archive Files file type is incremented by one. The number of macro-enabled attachments within an archive file are not counted.
- If one or more macros are detected within an embedded file, the parent file type is incremented by one. The number of macro-enabled attachments within an embedded file are not counted.

Virus Types Page

The **Email > Reporting > Virus Types** page provides an overview of the viruses that are sent to and from your network. The Virus Types page displays the viruses that have been detected by the virus scanning engines running on the Email Security appliances and are displayed on the Security Management appliance. Use this report to take action against a particular virus. For example, if you see that you are receiving a high volume of viruses known to be embedded in PDF files, you can create a filter action to quarantine messages with PDF attachments.

Note Outbreak Filters can quarantine these types of virus-infected messages with no user intervention.

If you run multiple virus scanning engines, the Virus Types page includes results from all enabled virus scanning engines. The name of the virus that appears on the page is determined by the virus scanning engines. If more than one scanning engine detects a virus, it is possible to have more than one entry for the same virus.

Table 17: Details on the Email Reporting Virus Types Page

Section	Description
Time Range (drop-down list)	A drop-down list that can range from a day to 90 days or a custom range. For more information on time ranges and customizing this for your needs, see the Choosing a Time Range for Reports , on page 26.
Top Incoming Virus Types Detected	This section displays a chart view of the viruses that have been sent to your network.
Top Outgoing Virus Types Detected	This section displays a chart view of the viruses that have been sent from your network.
Virus Types Detail	An interactive table that shows the details of each virus type.

Note To see which hosts sent virus-infected messages to your network, go to the Incoming Mail page, specify the same reporting period, and sort by virus positive. Similarly, to see which IP addresses have sent virus positive email within your network, view the Outgoing Senders page and sort by virus positive messages.

From the **Virus Types** page you can also generate a PDF or export raw data to a CSV file. For information on printing or exporting a file, see the Understanding the Email Reporting Pages, on page 40.

Note You can generate a scheduled report for the **Virus Types** page. See the Scheduling Email Reports, on page 83.

URL Filtering Page

- URL Filtering report modules are populated only if URL filtering is enabled.
- URL Filtering reports are available for incoming and outgoing messages.
- Only messages that are scanned by the URL filtering engine (either as part of anti-spam/outbreak filter scanning or through message/content filters) are included in these modules. However, not all of the results are necessarily specifically attributable to the URL Filtering feature.
- The Top URL Categories module includes all categories found in messages that have been scanned, whether or not they match a content or message filter.

- Each message can be associated with only one reputation level. For messages with multiple URLs, the statistics reflect the lowest reputation of any URL in the message.
- URLs in the global whitelist configured at Security Services > URL Filtering are not included in reports. URLs in whitelists used in individual filters are included in reports.
- Malicious URLs are URLs that Outbreak Filters have determined to have poor reputation. Neutral URLs are those that Outbreak Filters have determined to require click-time protection. Neutral URLs have therefore been rewritten to redirect them to the Cisco Web Security proxy.
- Results of URL category-based filters are reflected in content and message filter reports.
- Results of click-time URL evaluations by the Cisco Web Security proxy are not reflected in reports.

Web Interaction Tracking Page

- Web Interaction Tracking report modules are populated only if the Web Interaction Tracking feature is enabled on managed Email Security appliances.
- Web Interaction Tracking reports are available for incoming and outgoing messages.
- Only rewritten URLs (either by policy or Outbreak Filter) clicked by the end users are included in these modules.
- Web Interaction Tracking page includes the following reports:
  Top Rewritten Malicious URLs clicked by End Users. Click on a URL to view a detailed report that contains the following information:
  - A list of end users who clicked on the rewritten malicious URL.
  - Date and time at which the URL was clicked.
  - Whether the URL was rewritten by a policy or an outbreak filter.
  - Action taken (allow, block, or unknown) when the rewritten URL was clicked. Note that, if a URL was rewritten by outbreak filter and the final verdict is unavailable, the status is shown as unknown.

  Note Due to a limitation, status of all outbreak rewritten URLs are shown as unknown.

  Top End Users who clicked on Rewritten Malicious URLs

  Tracking Web Interaction Details. Includes the following information:
  - A list of all the rewritten URLs (malicious and unmalicious). Click on a URL to view a detailed report.
  - Action taken (allow, block, or unknown) when a rewritten URL was clicked.
    If the verdict of a URL (clean or malicious) was unknown at the time when the end user clicked it, the status is shown as unknown. This could be because the URL was under further scrutiny or the web server was down or not reachable at the time of the user click.
  - The number of times end users clicked on a rewritten URL. Click a number to view a list of all messages that contain the clicked URL.
- Note the following:

- If you have configured a content or message filter to deliver messages after rewriting malicious URLs and notify another user (for example, an administrator), the web interaction tracking data for the original recipient is incremented if the notified user clicks on the rewritten URLs.
- If you are sending a copy of quarantined messages containing rewritten URLs to a user other than the original recipient (for example, to an administrator) using the web interface, the web interaction tracking data for the original recipient is incremented if the other user clicks on the rewritten URLs.

Forged Email Detection Page

- The Forged Email Detection page includes the following reports:
  - **Top Forged Email Detection**. Displays the top ten users in the content dictionary that matched the forged From: header in the incoming messages.
  - **Forged Email Detection Details**. Displays a list of all the users in the content dictionary that matched the forged From: header in the incoming messages and for a given user, the number of messages matched.
- The Forged Email Detection reports are populated only if you are using the Forged Email Detection content filter or the forged-email-detection message filter.

Advanced Malware Protection (File Reputation and File Analysis) Reporting Pages

Requirements for File Analysis Report Details

(Cloud File Analysis) Ensure That the Management Appliance Can Reach the File Analysis Server

In order to obtain File Analysis report details, the appliance must be able to connect to the File Analysis server over port 443. See details in Firewall Information, on page 401

If your Cisco Content Security Management appliance does not have a direct connection to the internet, configure a proxy server for this traffic (See Upgrade and Update Settings, on page 302.) If you have already configured the appliance to use a proxy to obtain upgrades and service updates, the existing settings are used.

If you use an HTTPS proxy, the proxy must not decrypt the traffic; use a pass-through mechanism for communications with the File Analysis server. The proxy server must trust the certificate from the Fire Analysis server, but need not provide its own certificate to the File Analysis server.

(Cloud File Analysis) Configure the Management Appliance to Display Detailed File Analysis Results

In order to allow all content security appliances in your organization to display detailed results in the cloud about files sent for analysis from any Cisco Email Security appliance or Cisco Web Security appliance in your organization, you need to join all appliances to the same appliance group.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Select **Management Appliance > Centralized Services > Security Appliances**.

Step 3 Scroll to the File Analysis section.

Step 4 If your managed appliances are pointed at different File Analysis cloud servers, select the server from which to display result details.
Result details will not be available for files processed by any other cloud server.

Step 5 Enter the Analysis Group ID.

- If you enter the Group ID incorrectly or need to change it for any other reason, you must open a case with Cisco TAC.
- This change takes effect immediately; it does not require a Commit.
- It is suggested to use your CCOID for this value.
- This value is case-sensitive.
- This value must be identical on all appliances that will share data about files that are uploaded for analysis.
- An appliance can belong to only one group.
- You can add a machine to a group at any time, but you can add it only once.

Step 6 Click **Group Now**.

Step 7 Configure the same group on each Email Security appliance that will share data with this appliance.

What to Do Next

Related Topics

For Which Files Are Detailed File Analysis Results Visible in the Cloud? , on page 70

(On-Premises File Analysis) Activate the File Analysis Account

If you have deployed an on-premises (private cloud) Cisco AMP Threat Grid Appliance, you must activate the File Analysis account for your Cisco Content Security Management appliance in order to view report details available on the Threat Grid appliance. You generally only need to do this once.

Before You Begin

Ensure that you are receiving System alerts at Critical level.

Step 1 The first time you attempt to access File Analysis report details from the Threat Grid appliance, wait a few minutes and you will receive an alert that includes a link.
If you do not receive this alert, go to **Management Appliance > System Administration > Alerts** and click **View Top Alerts**.

Step 2 Click the link in the alert message.

Step 3 Activate your management appliance account.

Additional Requirements

For any additional requirements, see the Release Notes for your Security Management appliance release, available from
http://www.cisco.com/c/en/us/support/security/content-security-management-appliance/products-release-notes-list.html

Identifying Files by SHA-256 Hash

Because filenames can easily be changed, the appliance generates an identifier for each file using a Secure Hash Algorithm (SHA-256). If an appliance processes the same file with different names, all instances are recognized as the same SHA-256. If multiple appliances process the same file, all instances of the file have the same SHA-256 identifier.

In most reports, files are listed by their SHA-256 value (in an abbreviated format).

File Reputation and File Analysis Report Pages

Report	Description
Advanced Malware Protection	Shows file-based threats that were identified by the file reputation service. For files with changed verdicts, see the AMP Verdict updates report. Those verdicts are not reflected in the Advanced Malware Protection report. If a file extracted from a compressed or archived file is malicious, only the SHA value of the compressed or archived file is included in the Advanced Malware Protection report. **Note** From AsyncOS 9.6.5 onwards, Advanced Malware Protection report has been enhanced to display additional fields, graphs, and so on. The report displayed after the upgrade does not include the reporting data prior to the upgrade. To view the Advanced Malware Protection report prior to AsyncOS 9.6.5 upgrade, click on the hyperlink at the bottom of the page. The **Incoming Malware Files by Category** section shows the percentage of blacklisted file SHAs received from the AMP for Endpoints console that are categorised as **Custom Detection**. The threat name of a blacklisted file SHA obtained from AMP for Endpoints console is displayed as **Simple Custom Detection** in the Incoming Malware Threat Files section of the report. You can click on the link in the More Details section of the report to view the file trajectory details of a blacklisted file SHA in the AMP for Endpoints console You can view the **Low Risk** verdict details in the Incoming Files Handed by AMP section of the report.

Report	Description
Advanced Malware Protection File Analysis	Displays the time and verdict (or interim verdict) for each file sent for analysis. The appliance checks for analysis results every 30 minutes. To view more than 1000 File Analysis results, export the data as a .csv file. For deployments with an on-premises Cisco AMP Threat Grid Appliance: Files that are whitelisted on the AMP Threat Grid appliance show as "clean." For information about whitelisting, see the AMP Threat Grid documentation or online help. Drill down to view detailed analysis results, including the threat characteristics for each file. You can also search for additional information about an SHA, or click the link at the bottom of the file analysis details page to view additional details on the server that analyzed the file. To view details on the server that analyzed a file, see Requirements for File Analysis Report Details , on page 65. If a file extracted from a compressed or archived file is sent for analysis, only the SHA value of the extracted file is included in the File Analysis report. **Note** From AsyncOS 9.6.5 onwards, File Analysis report has been enhanced to display additional fields, graphs, and so on. The report displayed after the upgrade does not include the reporting data prior to the upgrade. To view the File Analysis report prior to AsyncOS 9.6.5 upgrade, click on the hyperlink at the bottom of the page.
Advanced Malware Protection Verdict Updates	Because Advanced Malware Protection is focused on targeted and zero-day threats, threat verdicts can change as aggregated data provides more information. The AMP Verdict Updates report lists the files processed by this appliance for which the verdict has changed since the message was received. For more information about this situation, see the documentation for your Email Security appliance. To view more than 1000 verdict updates, export the data as a .csv file. In the case of multiple verdict changes for a single SHA-256, this report shows only the latest verdict, not the verdict history. To view all affected messages for a particular SHA-256 within the maximum available time range (regardless of the time range selected for the report) click a SHA-256 link.

Viewing File Reputation Filtering Data in Other Reports

Data for file reputation and analysis is available in other reports where relevant. A Detected by Advanced Malware Protection column may be hidden by default in applicable reports. To display additional columns, click the Columns link at the bottom of the table.

For Which Files Are Detailed File Analysis Results Visible in the Cloud?

If you have deployed public-cloud File Analysis, you can view detailed results for all files uploaded from any managed appliance that has been added to the appliance group for File Analysis.

If you have added your management appliance to the group, you can view the list of managed appliances in the group by clicking the button on the **Management Appliance > Centralized Services > Security Appliances** page.

Appliances in the analysis group are identified by the File Analysis Client ID. To determine this identifier for a particular appliance, look in the following location:

Appliance	Location of File Analysis Client ID
Email Security appliance	Advanced Settings for File Analysis section on the **Security Services > File Reputation and Analysis** page.
Web Security appliance	Advanced Settings for File Analysis section on the **Security Services > Anti-Malware and Reputation** page.
Cisco Content Security Management appliance	At the bottom of the **Management Appliance > Centralized Services > Security Appliances** page.

Related Topics

- (Cloud File Analysis) Configure the Management Appliance to Display Detailed File Analysis Results , on page 66

Mailbox Auto Remediation

You can view the details of the mailbox remediation results using the Mailbox Auto Remediation report page. Use this report to view details such as:

- A list of recipients for whom the mailbox remediation was successful or unsuccessful
- Remedial actions taken on messages
- The filenames associated with a SHA-256 hash

The **Recipients for whom remediation was unsuccessful** field is updated in the following scenarios:

- The recipient is not a valid Office 365 user or the recipient does not belong to the Office 365 domain account configured on your appliance.

- The message containing the attachment is no longer available in the mailbox, for example, the end user deleted the message.
- There was a connectivity issue between your appliance and Office 365 services when the appliance was trying to perform the configured remedial action.

Click on a SHA-256 hash to view the related messages in Message Tracking.

TLS Connections Page

The **Email > Reporting > TLS Connections** page shows the overall usage of TLS connections for sent and received mail. The report also shows details for each domain sending mail using TLS connections.

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

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

Table 18: Details on the Email Reporting TLS Connections Page

Section	Description
Time Range (drop-down list)	A drop-down list that can range from a day to 90 days or a custom range. For more information on time ranges and customizing this for your needs, see the Choosing a Time Range for Reports , on page 26.
Incoming TLS Connections Graph	The graph displays a view of incoming TLS encrypted and unencrypted connections over the last hour, day, or week depending on the time frame that you have selected.
Incoming TLS Connections Summary	This table displays the total volume of incoming messages, the volume of encrypted and unencrypted messages, and the volume of successful and failed incoming TLS encrypted messages.
Incoming TLS Message Summary	This table displays a summary of the total volume of incoming messages.
Incoming TLS Connections Details	The table displays details for domains sending or receiving encrypted messages. For each domain, you can view the total number of connections, messages sent, and the number of TLS connections that were successful or failed. You can also view the percentage of successful and failed connections for each domain.
Outgoing TLS Connections Graph	The graph displays a view of outgoing TLS encrypted and unencrypted connections over the last hour, day, or week depending on the time frame that you have selected.

Section	Description
Outgoing TLS Connections Summary	This table displays the total volume of outgoing messages, the volume of encrypted and unencrypted messages, and the volume of successful and failed outgoing TLS encrypted messages.
Outgoing TLS Message Summary	This table shows the total volume of outgoing messages
Outgoing TLS Connections Details	The table displays details for domains sending or receiving encrypted messages. For each domain, you can view the total number of connections, messages sent, and the number of TLS connections that were successful or failed, and the last TLS status. You can also view the percentage of successful and failed connections for each domain.

Inbound SMTP Authentication Page

The Inbound SMTP Authentication page shows the use of client certificates and the SMTP AUTH command to authenticate SMTP sessions between the Email Security appliance and users' mail clients. If the appliance accepts the certificate or SMTP AUTH command, it will establish a TLS connection to the mail client, which the client will use to send a message. Since it is not possible for the appliance to track these attempts on a per-user basis, the report shows details on SMTP authentication based on the domain name and domain IP address.

Use this report to determine the following information:

- Overall, how many incoming connection use SMTP authentication?
- How many connections use a client certificated?
- How many connections use SMTP AUTH?
- What domains are failing to connect when attempting to use SMTP authentication?
- How many connections are successfully using the fall-back when SMTP authentication fails?

The Inbound SMTP Authentication page includes a graph for received connections, a graph for mail recipients who attempted an SMTP authentication connection, and a table with details on the attempts to authenticate connections.

The Received Connections graph shows the incoming connections from mail clients that attempt to authentication their connections using SMTP authentication over the time range you specify. The graph displays the total number of connections the appliance received, the number that did not attempt to authenticate using SMTP authentication, the number that failed and succeeded to authenticate the connection using a client certificate, and the number that failed and succeeded to authenticate using the SMTP AUTH command.

The Received Recipients graph displays the number of recipients whose mail clients attempted to authenticate their connections to the Email Security appliances to send messages using SMTP authentication. The graph also show the number of recipients whose connections were authenticated and the number of recipients whose connections were not authenticated.

The SMTP Authentication details table displays details for the domains whose users attempt to authenticate their connections to the Email Security appliance to send messages. For each domain, you can view the number

of connection attempts using a client certificate that were successful or failed, the number of connection attempts using the SMTP AUTH command that were successful or failed, and the number that fell back to the SMTP AUTH after their client certificate connection attempt failed. You can use the links at the top of the page to display this information by domain name or domain IP address.

Rate Limits Page

Rate Limiting by envelope sender allows you to limit the number of email message recipients per time interval from an individual sender, based on the mail-from address. The Rate Limits report shows you the senders who most egregiously exceed this limit.

Use this report to help you identify the following:

- Compromised user accounts that might be used to send spam in bulk.
- Out-of-control applications in your organization that use email for notifications, alerts, automated statements, etc.
- Sources of heavy email activity in your organization, for internal billing or resource-management purposes.
- Sources of large-volume inbound email traffic that might not otherwise be considered spam.

Note that other reports that include statistics for internal senders (such as Internal Users or Outgoing Senders) measure only the number of messages sent; they do not identify senders of a few messages to a large number of recipients.

The Top Offenders by Incident chart shows the envelope senders who most frequently attempted to send messages to more recipients than the configured limit. Each attempt is one incident. This chart aggregates incident counts from all listeners.

The Top Offenders by Rejected Recipients chart shows the envelope senders who sent messages to the largest number of recipients above the configured limit. This chart aggregates recipient counts from all listeners.

Rate Limiting settings, including “Rate Limit for Envelope Senders” settings, are configured on the Email Security appliance in Mail Policies > Mail Flow Policies. For more information on rate limiting, see the documentation or online help for your Email Security appliance.

Related Topics

- High Volume Mail , on page 61

Outbreak Filters Page

The **Email > Reporting > Outbreak Filters** page shows information about recent outbreaks and messages quarantined due to Outbreak Filters. You can use this page to monitor your defense against targeted virus, scam, and phishing attacks.

Use the Outbreak Filters page to answer the following types of questions:

- How many messages are quarantined and by which Outbreak Filters rule?
- How much lead time has the Outbreak Filters feature been providing for virus outbreaks?
- How do the local outbreaks compare to the global outbreaks?

- How long do messages stay in the Outbreak Quarantine?
- Which potentially malicious URLs are most frequently seen?

The Threats By Type section shows the different types of threat messages received by the appliance. The Threat Summary section shows a breakdown of the messages by Virus, Phish, and Scam.

The Past Year Outbreak Summary lists global as well as local outbreaks over the past year, allowing you to compare local network trends to global trends. The listing of global outbreaks is a superset of all outbreaks, both viral and non-viral, whereas local outbreaks are limited to virus outbreaks that have affected your appliance. Local outbreak data does not include non-viral threats. Global outbreak data represents all outbreaks detected by the Threat Operations Center which exceeded the currently configured threshold for the outbreak quarantine. Local outbreak data represents all virus outbreaks detected on this appliance which exceeded the currently configured threshold for the outbreak quarantine. The Total Local Protection Time is always based on the difference between when each virus outbreak was detected by the Threat Operations Center and the release of an anti-virus signature by a major vendor. Note that not every global outbreak affects your appliance. A value of "--" indicates either a protection time does not exist, or the signature times were not available from the anti-virus vendors (some vendors may not report signature times). This does not indicate a protection time of zero, rather it means that the information required to calculate the protection time is not available.

The Quarantined Messages section summarizes Outbreak Filters quarantining, and is a useful gauge of how many potential threat messages Outbreak Filters are catching. Quarantined messages are counted at time of release. Typically, messages will be quarantined before anti-virus and anti-spam rules are available. When released, they will be scanned by the anti-virus and anti-spam software and determined to be positive or clean. Because of the dynamic nature of Outbreak tracking, the rule under which a message is quarantined (and even the associated outbreak) may change while the message is in the quarantine. Counting the messages at the time of release (rather than the time of entry into the quarantine) avoids the confusion of having counts that increase and decrease.

The Threat Details listing displays information about specific outbreaks, including the threat category (virus, scam, or phishing), threat name, a description of the threat, and the number of messages identified. For virus outbreaks, the Past Year Virus Outbreaks include the Outbreak name and ID, time and date a virus outbreak was first seen globally, the protection time provided by Outbreak filters, and the number of quarantined messages. You can choose whether to view global or local outbreaks.

The First Seen Globally time is determined by the Threat Operations Center, based on data from the SenderBase, the world's largest email and web traffic monitoring network. The Protection Time is based on the difference between when each threat was detected by the Threat Operations Center and the release of an anti-virus signature by a major vendor.

A value of "--" indicates either a protection time does not exist, or the signature times were not available from the anti-virus vendors (some vendors may not report signature times). This does not indicate a protection time of zero. Rather, it means that the information required to calculate the protection time is not available.

Other modules on this page provide:

- The number of incoming messages processed by Outbreak Filters in the selected time period.

Non-viral threats include phishing emails, scams, and malware distribution using links to an external website.

- Severity of threats caught by Outbreak Filters.

Level 5 threats are severe in scope or impact, while Level 1 represents low threat risk. For descriptions of threat levels, see the online help or user guide for your Email Security appliance.

- Length of time messages spent in the Outbreak Quarantine.

This duration is determined by the time it takes the system to compile enough data about the potential threat to make a verdict on its safety. Messages with viral threats typically spend more time in the quarantine than those with non-viral threats, because they must wait for anti-virus program updates. The maximum retention time that you specify for each mail policy is also reflected.

- The URLs most frequently rewritten to redirect message recipients to the Cisco Web Security Proxy for click-time evaluation of the site if and when the recipient clicks a potentially malicious link in a message.

This list may include URLs that are not malicious, because if any URL in a message is deemed malicious, then all URLs in the message are rewritten.

Note In order to correctly populate the tables on the Outbreak Filters reporting page, the appliance must be able to communicate with the Cisco update servers specified in Management Appliance > System Administration > Update Settings.

For more information, see the Outbreak Filters chapter.

Reporting of Graymail

Graymail statistics are reflected in the following reports:

Report	Contains the Following Graymail Data
Mail Flow Summary page > Incoming tab	The number of incoming graymail messages under each graymail category (Marketing, Social, and Bulk) and the total number of graymail messages.
Mail Flow Details page > Outgoing Senders tab	The top graymail senders.
Mail Flow Details page > Incoming Mails tab	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.
User Mail Summary page > Top Users by Graymail	The top end users who receive graymail.
User Mail Summary page > User 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 users.

Reporting of Marketing Messages after Upgrade to AsyncOS 9.5

After upgrade to AsyncOS 9.5:

- The number of marketing messages is the 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.
- If the graymail feature is not enabled on managed Email Security appliances, marketing messages are counted as clean messages.

System Capacity Page

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

The System Capacity page can be used to determine the following information:

- Identify when Email Security appliances are exceeding recommended capacity; this enables you to determine when configuration optimization or additional appliances are needed.
- Identify historical trends in system behavior that point to upcoming capacity issues.
- For troubleshooting, identify which parts of the system are using the most resources.

Monitor your Email Security appliances to ensure that the capacity is appropriate to your message volumes. Over time, volume inevitably rises and appropriate monitoring ensures that additional capacity or configuration changes can be applied proactively. The most effective way to monitor system capacity is to track the overall volume, the messages in the work queue, and the incidents of Resource Conservation Mode.

- **Volume:**It is important to understand the “normal” message volume and the “usual” spikes in your environment. Track this data over time to measure volume growth. You can use the Incoming Mail and Outgoing Mail pages to track volume over time. For more information, see System Capacity – Incoming Mail, on page 77 and System Capacity – Outgoing Mail, on page 78.
- **Work Queue:** The work queue is designed to work as a “shock absorber”— absorbing and filtering spam attacks and processing unusual increases in non-spam messages. However, the work queue can also indicate a system under stress. Prolonged and frequent work queue backups may indicate a capacity problem. You can use the System Capacity – Workqueue page to track the activity in your work queue. For more information, see System Capacity – Workqueue, on page 77.
- **Resource Conservation Mode:** When an appliance becomes overloaded, it enters Resource Conservation Mode (RCM) and sends a CRITICAL system alert. This is designed to protect the device and allow it to process any backlog of messages. Your appliance should enter RCM infrequently and only during a very large or unusual increase in mail volume. Frequent RCM alerts may be an indication that the system is becoming overloaded. See Resource Conservation Activity , on page 78.

How to Interpret the Data You See on System Capacity Page

When choosing time ranges for viewing data on the System Capacity page, the following is important to remember:

- Day Report— The Day report queries the hour table and displays the exact number of queries that have been received by the appliance on an hourly basis over a 24 hour period. This information is gathered from the hour table. This is an exact number.

- Month Report— The Month report queries the day tables for the 30 or 31 days (dependent on the number of days in the month), giving you an exact report on the number of queries over 30 or 31 days. Again, this is an exact number.

The 'Maximum' value indicator on the System Capacity page is the highest value seen for the specified period. The 'Average' value is the average of all values for the specified period. The period of aggregation depends on the interval selected for that report. For example, you can choose to see the Average and Maximum values for each day if the chart is for a month period.

You can click the View Details link for a specific graph to view data for individual Email Security appliances and overall data for the appliances connected to the Security Management appliance.

System Capacity – Workqueue

The Workqueue page shows the average time a message spends in the work queue, excluding any time spent in the Spam quarantine or in a policy, virus, or outbreak quarantine. You can view time periods from an hour up to one month. This average can help in identifying both short term events delaying mail delivery and identify long term trends in the workload on the system.

Note If a message is released from the quarantine into the work queue, the "average time in work queue" metric ignores this time. This prevents double-counting and distorted statistics due to extended time spent in a quarantine.

The report also shows the volume of messages in the work queue over a specified time period, and it shows the maximum messages in the work queue over the same time period. The graphical representation of the maximum messages in the work queue also shows the work queue threshold level.

Occasional spikes in the Workqueue graphs are normal and expected. If the messages in the work queue remain higher than the configured threshold for a long duration, this may indicate a capacity issue. In this scenario, consider tuning the threshold level or review the system configuration.

To change the work queue threshold level, see Adjusting the Reference Threshold in System Health Graphs for Email Security Appliances, on page 337.

Tip When reviewing the work queue page, you may want to measure the frequency of work queue backups, and take note of work queue backups that exceed 10,000 messages.

System Capacity – Incoming Mail

The System Capacity – Incoming Mail page shows incoming connections, the total number of incoming messages, the average message size, and the total incoming message size. You can view the results for a day, week, month, or year. It is important to understand the trends of normal message volume and spikes in your environment. You can use the System Capacity – Incoming Mail page to track volume growth over time and plan for system capacity. You might also want to compare the incoming mail data with the sender profile data to view the trends in volumes of email messages that are sent from specific domains to your network.

Note An increased number of incoming connections may not necessarily affect system load.

System Capacity – Outgoing Mail

The System Capacity – Outgoing Mail page shows outgoing connections, the total number of outgoing messages, the average message size, and the total outgoing message size. You can view the results for a day, week, month, or year. It is important to understand the trends of normal message volume and spikes in your environment. You can use the System Capacity – Outgoing Mail page to track volume growth over time and plan for system capacity. You might also want to compare the outgoing mail data with the outgoing destinations data to view the trends in volumes of email messages that are sent from specific domains or IP addresses.

System Capacity – System Load

The system load report shows the following:

Overall CPU Usage

Email Security appliances are optimized to use idle CPU resources to improve message throughput. High CPU usage may not indicate a system capacity problem. If the high CPU usage is coupled with consistent, high-volume memory page swapping, you may have a capacity problem.

Note This graph also indicates a threshold for CPU usage that is a visual reference only. To adjust the position of this line, see Adjusting the Reference Threshold in System Health Graphs for Email Security Appliances, on page 337. You can configure your Email Security appliances to send you alerts that will suggest actions that you can take to address capacity issues.

This page also shows a graph that displays the amount of CPU used by different functions, including mail processing, spam and virus engines, reporting, and quarantines. The CPU-by-function graph is a good indicator of which areas of the product use the most resources on your system. If you need to optimize your appliance, this graph can help you determine which functions may need to be tuned or disabled.

Memory Page Swapping

The memory page swapping graph shows how frequently the system must page to disk, in kilobytes per second.

The system is designed to swap memory regularly, so some memory swapping is expected and is not an indication of problems with your appliance. Unless the system *consistently* swaps memory in high volumes, memory swapping is normal and expected behavior (especially on C170appliances). To improve performance, you may need to add Email Security appliances to your network or tune your configuration to ensure maximum throughput.

Note This graph also indicates a threshold for memory page swapping that is a visual reference only. To adjust the position of this line, see Adjusting the Reference Threshold in System Health Graphs for Email Security Appliances, on page 337. You can configure your Email Security appliances to send you alerts that will suggest actions that you can take to address capacity issues.

Resource Conservation Activity

The resource conservation activity graph shows the number of times the Email Security appliance entered Resource Conservation Mode (RCM). For example, if the graph shows n times, it means that the appliance has entered RCM n times and exited at least n-1 times.

Your appliances should enter RCM infrequently and only during a very large or unusual increase in mail volume. If the Resource Conservation Activity graph shows that your appliance is entering RCS frequently, it may be an indication that the system is becoming overloaded.

System Capacity – All

The **All** page consolidates all the previous system capacity reports onto a single page so you can view the relationship between the different reports. For example, you might see that the message queue is high at the same time that excessive memory swapping takes place. This might be an indication that you have a capacity problem. You may want to save this page as a PDF file to preserve a snapshot of system performance for later reference (or to share with support staff).

Threshold Indicator in System Capacity Graphs

In some graphs, a line indicates the default value that may indicate a possible problem if it is frequently or consistently crossed. To adjust this visual indicator, see Adjusting the Reference Threshold in System Health Graphs for Email Security Appliances, on page 337.

Reporting Data Availability Page

The **Email > Reporting > Reporting Data Availability** page allows you to view, update and sort data to provide real-time visibility into resource utilization and email traffic trouble spots.

All data resource utilization and email traffic trouble spots are shown from this page, including the data availability for the overall appliances that are managed by the Security Management appliance.

From this report page you can also view the data availability for a specific appliance and time range.

About Scheduled and On-Demand Email Reports

Types of Reports Available

Except as noted, the following types of Email Security reports are available as both scheduled and on-demand reports:

- Content Filters—This report includes up to 40 content filters. For additional information on what is included on this page, see the Content Filters Page, on page 144.
- DLP Incident Summary—For information on what is included on this page, see the DLP Incident Summary Page, on page 141.
- Delivery Status—The report page displays information about delivery problems to a specific recipient domain or Virtual Gateway address, page displays a list of the top 20, 50, or 100 recipient domains for messages delivered by the system within the last three hours. You can sort by latest host status, active recipients (the default), connections out, delivered recipients, soft bounced events, and hard bounced recipients by clicking the links in the column heading for each statistic. For more information on what the Delivery Status page does on the Email Security appliance, see the documentation or online help for your Email Security appliance.
- Domain-Based Executive Summary—This report is based on the Mail Flow Summary Page, on page 109, and is limited to a group of specified domains. For information on what is included, see the Domain-Based Executive Summary Report, on page 81.

- Executive Summary—This report is based on the information from the Mail Flow Summary Page, on page 109. For information on what is included, see the Domain-Based Executive Summary Report, on page 81.
- Mail Flow Details — For information on what in included on this page, see the Mail Flow Details Page, on page 126.
- User Mail Summary—For information on what is included on this page, see the User Mail Summary, on page 138.
- Outgoing Destinations—For information on what is included on this page, see the Outgoing Destinations Page, on page 131.
- Sender Groups —For information on what is included on this page, see the Sender Groups Page, on page 131.
- TLS Encrptions—For information on what is included on this page, see the TLS Encryption Page, on page 133.
- Virus Types—For information on what is included on this page, see the Virus Filtering Page, on page 123.

Time ranges

Depending on the report, these reports can be configured to include data for the previous day, previous seven days, previous month, previous calendar day (up to 250), or previous calendar month (up to 12). Alternatively, you can include data for a custom number of days (from 2 days to 100 days) or a custom number of months (from 2 months to 12 months).

Regardless of when you run a report, the data is returned from the previous time interval (hour, day, week, or month). For example, if you schedule a daily report to run at 1AM, the report will contain data from the previous day, midnight to midnight (00:00 to 23:59).

Languages and Locales

Note You can schedule a PDF report or export raw data as a CSV file with a specific locale for that individual report. The language drop-down menu on the Scheduled Reports page allows you to view or schedule a PDF report in the users current selected locale and language. See important information at Printing and Exporting Reporting Data , on page 31.

Storage of Archived Reports

For information on how long reports are stored for, and when archived reports are deleted from the system, see Viewing and Managing Archived Email Reports , on page 87.

Additional Report Types

Two special reports that can be generated in the **Email > Reporting** section on the Security Management appliance are:

Domain-Based Executive Summary Report

The Domain-Based Executive Summary report provides a synopsis of the incoming and outgoing message activity for one or more domains in your network. It is similar to the Executive Summary report, but it limits the report data to the messages sent to and from the domains that you specify. The outgoing mail summary shows data only when the domain in the PTR (pointer record) of the sending server matches a domain you specify. If multiple domains are specified, the appliance aggregates the data for all those domains into a single report.

To generate reports for a subdomain, you must add its parent domain as a second-level domain in the reporting system of the Email Security appliance and the Security Management appliance. For example, if you add example.com as a second-level domain, its subdomains, such as subdomain.example.com, are available for reporting. To add second-level domains, use **reportingconfig -> mailsetup -> tld** in the Email Security appliance CLI, and **reportingconfig -> domain -> tld** in the Security Management appliance CLI.

Unlike other scheduled reports, Domain-Based Executive Summary reports are not archived.

Domain-Based Executive Summary Reports and Messages Blocked by Sender Reputation Filtering

Because messages blocked by sender reputation filtering do not enter the work queue, AsyncOS does not process these messages to determine the domain destination. An algorithm estimates the number of rejected messages per domain. To determine the exact number of blocked messages per domain, you can delay HAT rejections on the Security Management appliance until the messages reach the recipient level (RCPT TO). This allows AsyncOS to collect recipient data from the incoming messages. You can delay rejections using **listenerconfig -> setup** command on the Email Security appliance. However, this option can impact system performance. For more information about delayed HAT rejections, see the documentation for your Email Security appliance.

Note To see Stopped by Reputation Filtering results in your Domain-Based Executive Summary report on the Security Management appliance, you must have **hat_reject_info** enabled on both the Email Security appliance and the Security Management appliance.To enable the **hat_reject_info** on the Security Management appliance, run the **reportingconfig > domain > hat_reject_info** command.

Managing Lists of Domains and Recipients for Domain-Based Executive Summary Reports

You can use a configuration file to manage the domains and recipients for a Domain-Based Executive Summary report. The configuration file is a text file that is stored in the configuration directory of the appliance. Each line in the file produces a separate report. This allows you to include a large number of domains and recipients in a single report, as well as define multiple domain reports in a single configuration file.

Each line of the configuration file includes a space-separated list of domain names and a space-separated list of email addresses for the report recipients. A comma separates the list of domain names from the list of email addresses. You can include subdomains by appending the subdomain name and a period at the beginning of the parent domain name, such as subdomain.example.com.

The following is a Single Report configuration file that generates three reports.

```
yourdomain.com sampledomain.com, admin@yourdomain.com
sampledomain.com, admin@yourdomain.com user@sampledomain.com
subdomain.example.com mail.example.com, user@example.com
```

Note You can use a configuration file and the settings defined for a single named report to generate multiple reports at the same time. For example, a company named Bigfish purchases two other companies, Redfish and Bluefish, and continues to maintain their domains. Bigfish creates a single Domain-Based Executive Summary report using a configuration file containing three lines corresponding to separate domain reports. When the appliance generates a Domain-Based Executive Summary report, an administrator for Bigfish receives a report on the Bigfish.com, Redfish.com, and Bluefish.com domains, while a Redfish administrator receives a report on the Redfish.com domain and a Bluefish administrator receives a report on the Bluefish.com domain.

You can upload a different configuration file to the appliance for each named report. You can also use the same configuration file for multiple reports. For example, you might create separate named reports that provide data about the same domains over different time periods. If you update a configuration file on your appliance, you do not have to update the report settings in the GUI unless you change the filename.

Creating Domain-Based Executive Summary Reports

Step 1 On the Security Management appliance, you can schedule the report or generate the report immediately.
To schedule the report:

a) [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.
b) Choose **Email > Reporting > Scheduled Reports**.
c) Click **Add Scheduled Report**.
To create an on-demand report:

- [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.
- Choose **Email > Reporting > Archived Reports.**
- Click **Generate Report Now**.

Step 2 From the **Report Type** drop-down list, choose **Domain-Based Executive Summary** report type.

Step 3 Specify the domains to include in the report and the email addresses for the report recipients. You can select one of the following options for generating the report:

- Generate report by specifying individual domains. Enter the domains for the report and the email addresses for the report recipients. Use commas to separate multiple entries. You can also use subdomains, such as subdomain.yourdomain.com. Specifying individual domains is recommended if you create reports for a small number of domains that are not expected to change frequently.
- Generate reports by uploading file. Import a configuration file that contains a list of the domains and recipient email addresses for the report. You can select a configuration file from the configuration directory on the appliance or upload one from your local computer. Using a configuration file is recommended if you create reports for a large number of domains that change frequently. For more information on configuration files for domain-based reports, see Managing Lists of Domains and Recipients for Domain-Based Executive Summary Reports , on page 81.

Note If you send reports to an external account (such as Yahoo! Mail or Gmail), you may need to add the reporting return address to the external account's whitelist to prevent report messages from being incorrectly classified as spam.

Step 4 In the Title text field, type the name of the title for the report.
AsyncOS does not verify the uniqueness of report names. To avoid confusion, do not create multiple reports with the same name.

Step 5 In the Outgoing Domain section, choose the domain type for the outgoing mail summary. Choices are: By Server or By Email Address.

Step 6 From the Time Range to Include drop-down list, select a time range for the report data.

Step 7 In the Format section, choose the format of the report.
Choices include:

- PDF. Create a formatted PDF document for delivery, archival, or both. You can view the report as a PDF file immediately by clicking Preview PDF Report.
- CSV. Create an ASCII text file that contains raw data as comma-separated values. Each CSV file may contain up to 100 rows. If a report contains more than one type of table, a separate CSV file is created for each table.

Step 8 From the Schedule section, choose a schedule for generating the report.
Choices include: Daily, Weekly (drop-down list for day of week included), or monthly.

Step 9 (Optional) Upload a custom logo for the report. The logo appears at the top of the report.

- The logo should be a .jpg, .gif, or .png file that is at most 550 x 50 pixels.
- If a logo file is not supplied, the default Cisco logo is used.

Step 10 Select a language for this report. For generating PDFs in Asian languages, see important information at Printing and Exporting Reporting Data , on page 31.

Step 11 Click **Submit** to submit your changes on the page, then click **Commit Changes** to commit your changes.

Executive Summary Report

The Executive Summary Report is a high-level overview of the incoming and outgoing email message activity from your Email Security appliances. that can be viewed on the Security Management appliance.

This report page summarizes what you can view on the Mail Flow Summary Page, on page 109. For more information on the Email Reporting Overview page, see Mail Flow Summary Page, on page 109.

Scheduled Reports Page

- Scheduling Email Reports, on page 83

Scheduling Email Reports

You can schedule any of the reports listed in About Scheduled and On-Demand Email Reports , on page 79.

To manage report scheduling, see the following:

Adding Scheduled Reports

To add a scheduled email report, use the following steps:

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Scheduled Reports**.

Step 3 Click **Add Scheduled Report**.

Step 4 Choose your report type.
For descriptions of the report types, see About Scheduled and On-Demand Email Reports , on page 79.

Note - For information about the settings for a Domain-Based Executive Summary report, see Domain-Based Executive Summary Report, on page 81.

- Available options for scheduled reports differ by report type. Options described in the remainder of this procedure do not necessarily apply to all reports.

Step 5 In the **Title** field, type the title of your report.
To avoid creating multiple reports with the same name, we recommend using a descriptive title.

Step 6 Choose the time range for the report from the **Time Range to Include** drop-down menu.

Step 7 Choose the format for the generated report.
The default format is PDF. Most reports also allow you to save raw data as a CSV file.

Step 8 Depending on the report, for Number of Rows, choose the amount of data to include.

Step 9 Depending on the report, choose the column by which to sort the report.

Step 10 From the **Schedule** area, select the radio button next to the day, week, or month for your scheduled report. Additionally, include the time that you want the report scheduled for. Time increments are based on midnight to midnight (00:00 to 23:59).

Step 11 In the **Email** text field, type in the email address where the generated report will be sent.
If you do not specify an email recipient, the system will still archive the reports.

You can add as many recipients for reports as you want, including zero recipients. If you need to send the reports to a large number of addresses, however, you may want to create a mailing list instead of listing the recipients individually.

Step 12 Choose a language for the report.
For Asian languages, see important information at Printing and Exporting Reporting Data , on page 31.

Step 13 Click **Submit**.

Editing Scheduled Reports

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Scheduled Reports**.

Step 3 Click the report name link in the Report Title column that you want to modify.

Step 4 Modify the report settings.

Step 5 Submit and commit your changes.

Discontinuing Scheduled Reports

To prevent future instances of scheduled reports from being generated, perform the following steps:

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Scheduled Reports**.

Step 3 Select the check boxes corresponding to the reports that you want to discontinue generating. To remove all scheduled reports, select the **All** check box.

Step 4 Click **Delete**.

Note Any archived versions of deleted reports are *not* automatically deleted. To delete previously-generated reports, see Deleting Archived Reports, on page 87.

Generating Email Reports On Demand

In addition to the reports that you can view (and generate PDFs for) using the interactive report pages described in Understanding the Email Reporting Pages, on page 106, you can save PDFs or raw-data CSV files for the reports listed in About Scheduled and On-Demand Email Reports , on page 79 at any time, for the time frame that you specify.

To generate an on-demand report perform the following:

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Archived Reports**.

Step 3 Click **Generate Report Now**.

Step 4 Choose a report type.

For descriptions of the report types, see About Scheduled and On-Demand Email Reports , on page 79.

Step 5 In the Title text field, type the name of the title for the report.
AsyncOS does not verify the uniqueness of report names. To avoid confusion, do not create multiple reports with the same name.

Note For information about the settings for a Domain-Based Executive Summary report, see Domain-Based Executive Summary Report, on page 81.

Available options for scheduled reports differ by report type. Options described in the remainder of this procedure do not necessarily apply to all reports.

Step 6 From the Time Range to Include drop-down list, select a time range for the report data.
Note the custom time range option.

Step 7 In the Format section, choose the format of the report.
Choices include:

- PDF. Create a formatted PDF document for delivery, archival, or both. You can view the report as a PDF file immediately by clicking Preview PDF Report.
- CSV. Create an ASCII text file that contains raw data as comma-separated values. Each CSV file may contain up to 100 rows. If a report contains more than one type of table, a separate CSV file is created for each table.

Step 8 Select the appliances or appliance groups for which you want to run the report. If you have not created any appliance groups, this option does not appear.

Step 9 From the Delivery Option section, choose the following:

- Archive the report by checking the **Archive Report** checkbox.

By choosing this, the report will be listed on the Archived Reports page.

Note Domain-Based Executive Summary reports cannot be archived.

- Email the report, by checking the **Email now to recipients** checkbox.

In the text field, type in the recipient email addresses for the report.

Step 10 Select a language for this report. For generating PDFs in Asian languages, see important information at Printing and Exporting Reporting Data , on page 31.

Step 11 Click **Deliver This Report** to generate the report.

Archived Email Reports Page

Viewing and Managing Archived Email Reports

Scheduled and on-demand reports are archived for a period of time.

The Security Management appliance retains the most recent reports that it generates, up to 30 instances of each scheduled report, up to 1000 total versions for all reports. The limit of 30 instances applies to each scheduled report with the same name and time range.

Archived reports are deleted automatically. As new reports are added, older reports are removed to keep the number at 1000.

Archived reports are stored in the /periodic_reports directory on the appliance. (See IP Interfaces and Accessing the Appliance, on page 389 for more information.)

Accessing Archived Reports

The **Email > Reporting > Archived Reports** page lists scheduled and on-demand reports that you have chosen to archive which have been generated and not yet purged.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Archived Reports**.

Step 3 To locate a particular report if the list is long, filter the list by choosing a report type from the **Show** menu, or click a column heading to sort by that column.

Step 4 Click the Report Title to view that report.

Deleting Archived Reports

Reports are automatically deleted from the system according to the rules outlined in Viewing and Managing Archived Email Reports , on page 87. However, you can manually delete unneeded reports.

To manually delete Archived reports, perform the following:

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Archived Reports**.
The Archived reports that are available are displayed.

Step 3 Select the checkbox for one or more reports to delete.

Step 4 Click **Delete**.

Step 5 To prevent future instances of scheduled reports from being generated, see Discontinuing Scheduled Reports , on page 85.

Troubleshooting Email Reports

See also Troubleshooting All Reports , on page 33.

Outbreak Filters Reports Do Not Show Information Correctly

Problem

Outbreak Filters reports do not show threat information correctly.

Solution

Verify that the appliance can communicate with the Cisco update servers specified in Management Appliance > System Administration > Update Settings.

Message Tracking Results Do Not Match Report Results After Clicking a Link in a Report

Problem

Message tracking results when drilling down from reports do not match expected results.

Solution

This can occur if reporting and tracking are not consistently and simultaneously enabled and functioning properly, or are not consistently and simultaneously either centralized or stored locally on each Email Security appliance. Data for each feature (reporting, tracking) is captured only while that feature is enabled.

Related Topics

- Checking Message Tracking Data Availability , on page 163

Advanced Malware Protection Verdict Updates Report Results Differ

Problem

A Web Security appliance and an Email Security appliance sent the same file for analysis, and the AMP Verdict Updates reports for Web and Email show different verdicts for that file.

Solution

This situation is temporary. Results will match once all verdict updates have been downloaded. Allow up to 30 minutes for this to occur.

Issues Viewing File Analysis Report Details

File Analysis Report Details Are Not Available

Problem

File Analysis report details are not available.

Solution

See Requirements for File Analysis Report Details , on page 65.

Error When Viewing File Analysis Report Details

Problem

No cloud server configuration is available error appears when you attempt to view File Analysis report details.

Solution

Go to **Management Appliance > Centralized Services > Security Appliances** and add at least one Email Security appliance that has the File Analysis feature enabled.

Error When Viewing File Analysis Report Details with Private Cloud Cisco AMP Threat Grid Appliance

Problem

You see an API key, registration, or activation error when attempting to view File Analysis report details.

Solution

If you are using a private cloud (on-premises) Cisco AMP Threat Grid appliance for file analysis, see (On-Premises File Analysis) Activate the File Analysis Account , on page 66.

If your Threat Grid appliance hostname changes, you must repeat the process in the referenced procedure.

Logging of File Analysis-Related Errors

Registration and other File Analysis-related errors are logged in the GUI logs.

Total Graymail or Marketing Messages Appears To Be Incorrect

Problem

The count of Marketing, Social and Bulk mail exceeds the total number of graymail messages.

Solution

The total number of Marketing Messages includes marketing messages received both before and after upgrade to AsyncOS 9.5, but the total number of graymail messages includes only messages received after upgrade. See Reporting of Marketing Messages after Upgrade to AsyncOS 9.5 , on page 75.

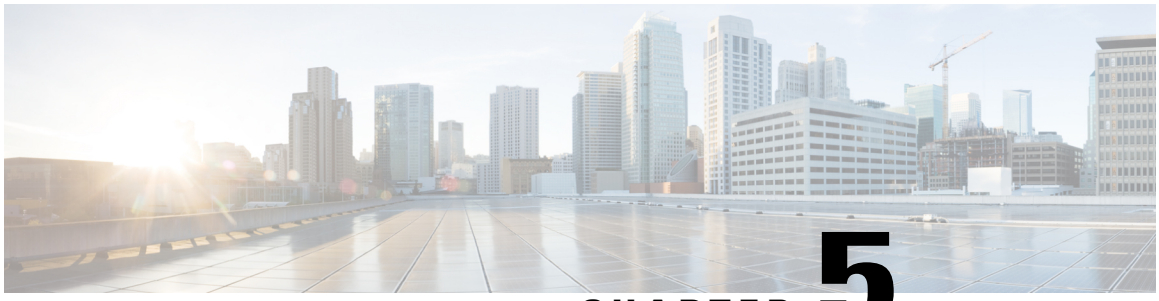

CHAPTER 5

Working With Reports on the Cloud Email Security Management Console

This chapter contains the following sections:

Ways to View Reporting Data

The following table shows the different ways to view reporting data:

Table 19: Ways To View Reporting Data

To	See
View and customize web-based interactive report pages	• Using the Interactive Report Pages, on page 93 • Customizing Your View of Report Data , on page 94 • Using Centralized Email Security Reporting on the Legacy Web Interface, on page 35
Automatically generate recurring CSV reports	Scheduling Email Reports, on page 83

To	See
Generate a CSV report on demand	Generating Email Reports On Demand , on page 85
Export raw data as a CSV (Comma-separated values) file	Printing and Exporting Reporting Data , on page 31 Exporting Report Data as a Comma Separated Values (CSV) File , on page 32
Email report information to yourself and other people	Scheduling Email Reports, on page 83 Generating Email Reports On Demand , on page 85
Find information about specific transactions	Viewing Details of Messages or Transactions Included in Reports , on page 30

Note For differences between logging and reporting, see Logging Versus Reporting, on page 347.

How the Security Management Appliance Gathers Data for Reports

The Security Management appliance pulls data for all reports from all managed appliances approximately every 15 minutes and aggregates the data from these appliances. Depending on your appliance, it may take awhile for a particular message to be included in the reporting data on the Security Management appliance. Check the **System Status** page for information on your data.

Reporting Data includes transactions involving both IPv4 and IPv6.

Note When gathering data for reports, the Security Management appliance applies the timestamp from the information that was set when you configured the time settings on the Security Management appliance. For information on setting the time on your Security Management appliance, see the Configuring the System Time, on page 324.

How Reporting Data is Stored

All of the appliances store reporting data. The following table shows what time periods that each appliance stores data.

Table 20: Reporting Data Storage on the Email Security Appliances

	Minute	Hourly	Daily	Weekly	Monthly	Yearly
Local Reporting on Email Security appliance	•	•	•	•	•	

	Minute	Hourly	Daily	Weekly	Monthly	Yearly
Centralized Reporting on Email Security appliance	•	•	•	•		
Security Management appliance		•	•	•	•	•

About Reporting and Upgrades

New reporting features may not apply to transactions that occurred before upgrade, because the required data may not have been retained for those transactions. For possible limitations related to reporting data and upgrades, see the Release Notes for your release.

Using the Interactive Report Pages

You can view the email reports for the Email Security appliance using the Reports drop-down as shown in the following figure:

Figure 3: Reports Drop-down

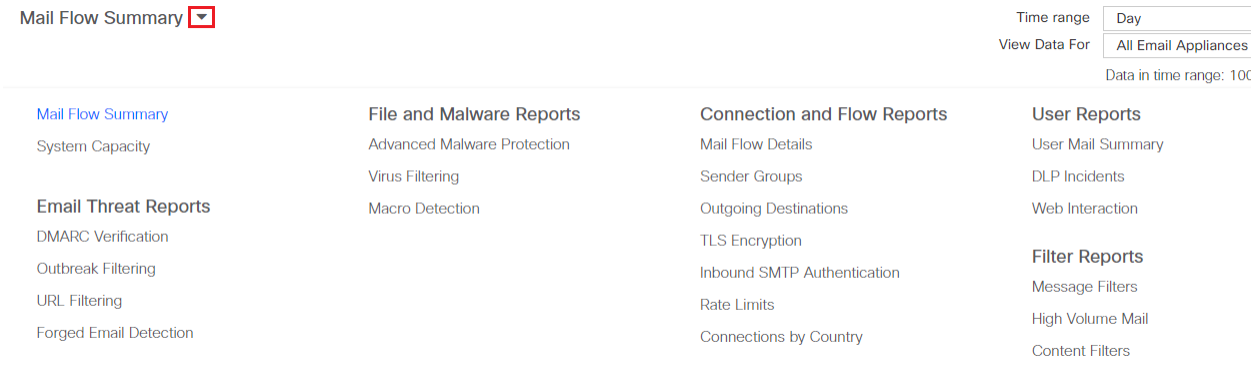

The Mail Flow Summary report page is the landing page (the page displayed after login).

Email reports are categorized into the following report pages under the reporting drop-down:

- Email Threat Reports
- File and Malware Reports
- Connection and Flow Reports
- User Reports
- Filter Reports

Customizing Your View of Report Data

You can customize your view while viewing the report data in the web interface.

To	Do This
Specify a time range	See Choosing a Time Range for Reports , on page 26.
View data per appliance or reporting group	See Viewing Reporting Data for an Appliance or Reporting Group, on page 26
Customize tables	See Customizing Tables on Report Pages , on page 95
Customize views	See Customizing Views on Report Pages, on page 95
Using Counters to Filter Data on Trend Graph	See Using Counters to Filter Data on the Trend Graphs, on page 96
Specify report-related preferences	See Setting Preferences , on page 346
Search for specific information or a subset of data to view	• For Email reports, see Searching and the Interactive Email Report Pages , on page 39. • Some tables include links (in blue text) to details for aggregated data. For more information, see Viewing Details of Messages or Transactions Included in Reports , on page 30.

Note All customization features are not available for every report.

Viewing Reporting Data for an Appliance or Reporting Group

For Mail Flow Summary and System Capacity reports for Email, you can view data from all appliances, or from any one centrally-managed appliance.

For Email reports, if you have created groups of Email Security appliances as described in Creating Email Reporting Groups, on page 37, you can view the data for each reporting group.

To specify the view, select an appliance or group from the **View Data For** list on supported pages.

If you are viewing report data on the Cloud Email Security Management Console to which you have recently taken backup from another Security Management appliance, you must first add (but do not establish a connection to) each appliance in > **Management Appliance** > **Centralized Services** > **Security Appliances**.

Choosing a Time Range for Reports

Most predefined report pages allow you to choose a Time Range for the data to include. The time range that you select is used for all of the report pages until you select a different value in the Time Range menu.

Available Time Range options differ by appliance and differ for Email reporting on the Security Management appliance:

Note Time ranges on report pages are displayed as a Greenwich Mean Time (GMT) offset. For example, Pacific time is GMT + 7 hours (GMT + 07:00).

Note All reports display date and time information based on the systems configured time zone, shown as a Greenwich Mean Time (GMT) offset. However, data exports display the time in GMT to accommodate multiple systems in multiple time zones around the world.

Tip You can specify a default time range that will always display each time you log in. For information, see Setting Preferences , on page 346.

Customizing Views on Report Pages

Most report pages allow you to choose between graphical view, tabular view or combined view. The view that you select is used to show the data on the report pages.

Table 21: Customizing Views on Report Pages

To	Do This
Show data in graph view.	Click  to view data in graphical format.
Show data in table view.	Click  to view data in tabular format.
View details about a table entry, where available	Click a blue entry in the table.
Show data in combined view.	Click  to view data in graphical and tabular format.

Customizing Tables on Report Pages

You can view, customize and sort information on the interactive tables within the report pages. The view that you select is used to show the data on the report pages.

Table 22: Customizing Tables on Report Pages

• Show additional columns • Hide visible columns • Determine available columns for a table	1 Click  . 2 Select the columns to display, and click **Close**.	For most tables, some columns are hidden by default. Each report page offers different columns. See the table column descriptions for the respective tables.
Sort the table by the heading of your choice.	Click a column heading.	-
View details about a table entry, where available	Click a blue entry in the table	See also Viewing Details of Messages or Transactions Included in Reports , on page 30.
View details of additional rows.	Scroll down on a table to display details of additional rows.	-

Using Counters to Filter Data on the Trend Graphs

You can filter data based on the required time range and available counters on a trend graph.

The time range that you select in the Time Range drop-down, is used for a trend graph until you select a different value.

A counters on a trend graph of the Mail Flow Summary report page is used to view data specific to different filters. Click on an available counter to filter the data.

Viewing Details of Messages or Transactions Included in Reports

Step 1 Click any blue number in a table on a report page.
(Not all tables have these links.)

The messages or transactions included in that number are displayed in Message Tracking.

Step 2 Scroll down to see the list of messages or transactions.

What to Do Next

- Tracking Messages, on page 159

Improving Performance of Email Reports

If the performance of aggregated reporting decreases due to a large number of unique entries over the course of a month, use reporting filters to restrict the aggregation of data in reports that cover the previous year (Last Year reports). These filters can restrict detailed, individual IP, domain, or user data in reports. Overview reports and summary information remain available for all reports.

You can enable one or more of the reporting filters using the **reportingconfig > filters** menu in the CLI. The changes must be committed to take effect.

- **IP Connection Level Detail**. Enabling this filter prevents the Security Management appliance from recording information about individual IP addresses. This filter is appropriate for systems that process a large number of incoming IP addresses due to attacks.

  This filter affects the following Last Year reports:

  - Sender Profile for Incoming Mail
  - IP Addresses for Incoming Mail
  - IP Addresses for Outgoing Senders

- **User Detail**. Enabling this filter prevents the Security Management appliance from recording information about individual users sending and receiving mail and the content filters that are applied to the users' mail. This filter is appropriate for appliances that process mail for millions of internal users or if the system does not validate recipient addresses.

  This filter affects the following Last Year reports:

  - Internal Users
  - Internal User Details
  - IP Addresses for Outgoing Senders
  - Content Filters

- **Mail Traffic Detail**. Enabling this filter prevents the Security Management appliance from recording information about individual domains and networks that the appliances monitor. This filter is appropriate when the number of valid incoming or outgoing domains is measured in the tens of millions.

  This filter affects the following Last Year reports:

  - Domains for Incoming Mail
  - Sender Profile for Incoming Mail
  - Internal User Details
  - Domains for Outgoing Senders

Note To view up-to-the-minute reporting data for the preceding hour, you must log in to an individual appliance and view the data there.

Printing and Exporting Reporting Data

Table 23: Printing and Exporting Reporting Data

To Get This	CSV	Do This	Notes
Raw data See also Exporting Report Data as a Comma Separated Values (CSV) File , on page 32	•	Click the Export link below the chart or table.	The CSV file contains all applicable data, including the data visible in the chart or table.
	•	Create a scheduled or on-demand report. See: • Generating Email Reports On Demand , on page 85 • Scheduling Email Reports, on page 83	Each CSV file may contain up to 100 rows. If a report contains more than one table, a separate CSV file is created for each table. Some extended reports are not available in CSV format.

Exporting Report Data as a Comma Separated Values (CSV) File

You can export raw data to a comma-separated values (CSV) file, which you can access and manipulate using database applications such as Microsoft Excel. For different ways to export data, see Printing and Exporting Reporting Data , on page 31.

Because CSV exports include only raw data, exported data from a web-based report page may not include calculated data such as percentages, even if that data appears in the web-based report.

For email message tracking and reporting data, the exported CSV data will display all data in GMT regardless of what is set on the Security Management appliance. This simplifies using data independently from the appliance, particularly when referencing data from appliances in multiple time zones.

The following example is an entry from a raw data export of the Anti-Malware category report, where Pacific Daylight Time (PDT) is displayed as GMT - 7 hours:

Begin Timestamp, End Timestamp, Begin Date, End Date, Name, Transactions Monitored, Transactions Blocked, Transactions Detected

1159772400.0, 1159858799.0, 2006-10-02 07:00 GMT, 2006-10-03 06:59 GMT, Adware, 525, 2100, 2625

Table 24: Viewing Raw Data Entries

Category Header	Value	Description
Begin Timestamp	1159772400.0	Query start time in number of seconds from epoch.

Category Header	Value	Description
End Timestamp	1159858799.0	Query end time in number of seconds from epoch.
Begin Date	2006-10-02 07:00 GMT	Date the query began.
End Date	2006-10-03 06:59 GMT	Date the query ended.
Name	Adware	Name of the malware category.
Transactions Monitored	525	Number of transactions monitored.
Transactions Blocked	2100	Number of transactions blocked.
Transactions Detected	2625	Total number of transactions: Number of transactions detected + Number of transactions blocked.

Note Category headers are different for each type of report.If you export localized CSV data, the headings may not be rendered properly in some browsers. This occurs because some browsers may not use the proper character set for the localized text. To work around this problem, you can save the file to your local machine, and open the file on any web browser using **File > Open**. When you open the file, select the character set to display the localized text.

Troubleshooting All Reports

Unable to View Report Data on Backup Security Management Appliance

Problem

You are unable to select a single Email Security appliance for which to view report data. The **View Data For** option does not appear on the reporting page.

Solution

See also Availability of Services During Backups , on page 292.

Reporting Is Disabled

Problem

Canceling a backup in progress can disable reporting.

Solution

Reporting functionality will be restored after a backup is completed.

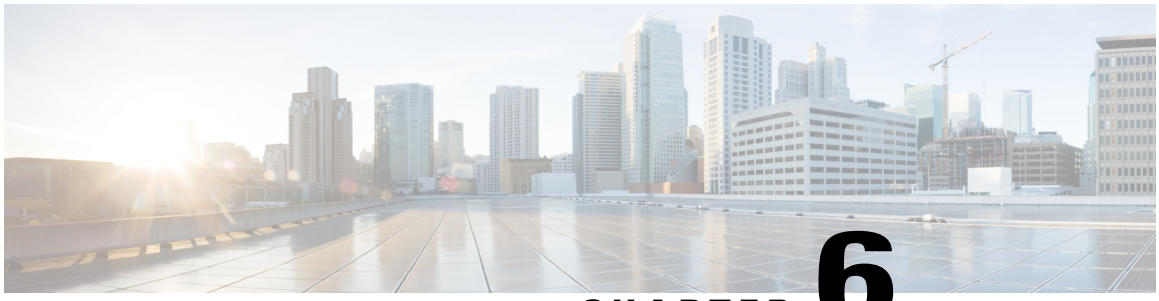

CHAPTER 6

Using Centralized Email Security Reporting on the Cloud Email Security Management Console

This chapter contains the following sections:

Centralized Email Reporting Overview

Your Cisco Content Security Management appliance shows aggregated information from individual or multiple Email Security appliances so that you can monitor your email traffic patterns and security risks. You can run reports in real-time to view an interactive display of system activity over a specific period of time, or you can schedule reports and run them at regular intervals. Reporting functionality also allows you to export raw data to a file.

This feature centralizes the reports listed under the Monitor menu of the Email Security appliance.

The Centralized Email Reporting feature not only generates high-level reports, allowing you to understand what is happening on their network, but it also allows you to drill down and see traffic details for a particular domain, user, or category.

The Centralized Tracking feature allows you to track email messages that traverse multiple Email Security appliances.

Note The Email Security appliance only stores data if local reporting is used. If centralized reporting is enabled for the Email Security appliance then the Email Security appliance does NOT retain any reporting data except for System Capacity and System Status. If Centralized Email Reporting is not enabled, the only reports that are generated are System Status and System Capacity.

For more information about availability of report data during and after the transition to centralized reporting, see the "Centralized Reporting Mode" section of the documentation or online help for your Email Security appliance.

Setting Up Centralized Email Reporting

To set up centralized email reporting, complete the following procedures in order:

Note If reporting and tracking are not consistently and simultaneously enabled and functioning properly, or are not consistently and simultaneously either centralized or stored locally on each Email Security appliance, then the message tracking results when drilling down from reports will not match expected results. This is because the data for each feature (reporting, tracking) is captured only while that feature is enabled.

Enabling Centralized Email Reporting on the Security Management Appliance

Before You Begin

- All Email Security appliances should be configured and working as expected before you enable centralized reporting.
- Before enabling centralized email reporting, ensure that sufficient disk space is allocated to that service. See the Managing Disk Space , on page 333.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance** > **Centralized Services** > **Centralized Reporting**.

Step 3 Click **Enable**.

Step 4 If you are enabling centralized email reporting for the first time after running the System Setup Wizard, review the end user license agreement, and click **Accept**.

Step 5 Submit and commit your changes.

Note If you have enabled email reporting on the appliance, and there is no disk space allocated for this action, centralized email reporting will not work until disk space is allocated. As long as the quota you are setting the Email Reporting and Tracking to is larger than the currently used disk space, you will not lose any reporting and tracking data. See the Managing Disk Space , on page 333 section, for more information.

Adding the Centralized Email Reporting Service to Each Managed Email Security Appliance

The steps you follow depend on whether or not you have already added the appliance while configuring another centralized management feature.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Security Appliances**.

Step 3 If you have already added the Email Security appliance to the list on this page:

a) Click the name of an Email Security appliance.

b) Select the **Centralized Reporting** service.

Step 4 If you have not yet added Email Security appliances:

a) Click Add Email Appliance.

b) In the Appliance Name and IP Address text fields, type the appliance name and the IP address for the Management interface of the Security Management appliance.

Note If you enter A DNS name in the IP Address text field, it will be immediately resolved to an IP address when you click **Submit**.

c) The Centralized Reporting service is pre-selected.

d) Click **Establish Connection.**

e) Enter the user name and password for an administrator account on the appliance to be managed, then click **Establish Connection.**.

Note You enter the login credentials to pass a public SSH key for file transfers from the Security Management appliance to the remote appliance. The login credentials are not stored on the Security Management appliance.

f) Wait for the Success message to appear above the table on the page.

g) Click **Test Connection.**.

h) Read test results above the table.

Step 5 Click **Submit**.

Step 6 Repeat this procedure for each Email Security appliance for which you want to enable Centralized Reporting.

Step 7 Commit your changes.

Enabling Centralized Email Reporting on Email Security Appliances

You must enable centralized email reporting on each managed Email Security appliance appliance.

For instructions, see the “Configuring an Email Security Appliance to Use Centralized Reporting” section of the documentation or online help for your Email Security appliance.

Creating Email Reporting Groups

You can create groups of Email Security appliances for which to view reporting data from the Security Management appliance.

A group can include one or more appliances, and an appliance may belong to more than one group.

Before You Begin

Make sure centralized reporting is enabled for each appliance. See Adding the Centralized Email Reporting Service to Each Managed Email Security Appliance , on page 103.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Centralized Reporting**.

Step 3 Click **Add Group** .

Step 4 Enter a unique name for the group.
The Email Security appliance list displays the Email Security appliances that you added to the Security Management appliance. Select the appliances that you want to add to the group.

The maximum number of groups that can be added is smaller than or equal to the maximum number of email appliances that can be connected.

Note If you added an Email Security appliance to the Security Management appliance, but you do not see it in the list, edit the configuration of the Email Security appliance so that the Security Management appliance is collecting reporting data from it.

Step 5 Click **Add** to add the appliances to the Group Members list.

Step 6 Submit and commit your changes.

Working with Email Report Data

- For options for accessing and viewing report data, see Ways to View Reporting Data , on page 91.
- To customize your view of report data, see Customizing Your View of Report Data , on page 94.
- To print or export report information, see Printing and Exporting Reporting Data , on page 31.
- To understand the various interactive report pages, see Using the Interactive Report Pages, on page 93.
- To generate a report on demand, see Generating Email Reports On Demand , on page 85.
- To schedule reports to run automatically at intervals and times that you specify, see Scheduling Email Reports, on page 83.
- To view archived on-demand and scheduled reports, see Viewing and Managing Archived Email Reports , on page 87.
- For background information, How the Security Management Appliance Gathers Data for Reports, on page 24.

- To improve performance when working with large amounts of data, see Improving Performance of Email Reports , on page 30.
- To get details about an entity or number that appears as a blue link in a chart or table, click the entity or number.

  For example, if your permissions allow you to do so, you can use this feature to view details about messages that violate Content Filtering or Data Loss Prevention policies. This performs the relevant search in Message Tracking. Scroll down to view results.

Searching and the Interactive Email Report Pages

Many of the interactive email reporting pages include a '**Search For:**' drop-down menu at the bottom of the page.

From the drop-down menu, you can search for several types of criteria, including the following:

- IP address
- Domain
- Network owner
- Internal user
- Destination domain
- Internal sender domain
- Internal sender IP address
- Incoming TLS domain
- Outgoing TLS domain
- SHA-256

For most searches, choose whether to exactly match the search text or look for items starting with the entered text (for example, starts with "ex" will match "example.com").

For IPv4 searches, the entered text is always interpreted as the beginning of up to four IP octets in dotted decimal format. For example, '17.*' will search in the range 17.0.0.0 through 17.255.255.255, so it will match 17.0.0.1 but not 172.0.0.1. For an exact match search, enter all four octets. IP address searches also support Classless Inter-Domain Routing (CIDR) format (17.16.0.0/12).

For IPv6 searches, you can enter addresses using the formats in the following examples:

- 2001:db8:2004:4202::0-2001:db8:2004:4202::ff
- 2001:db8:2004:4202::
- 2001:db8:2004:4202::23
- 2001:db8:2004:4202::/64

Understanding the Email Reporting Pages

Note This list represents the reports available in the latest supported release of AsyncOS for Email Security appliances. If your Email Security appliances are running earlier releases of AsyncOS, not all of these reports are available.

Table 25: Email Reporting Drop-down Options

Email Reporting Menu	Action
Mail Flow Summary Page	The Mail Flow Summary report page provides a synopsis of the activity on your Email Security appliances. It includes graphs and summary tables for the incoming and outgoing messages. For more information, see the Mail Flow Summary Page, on page 109.
Email Threat Reports	
DMARC Verification Page	The DMARC Verification report page shows the top sender domains that failed Domain-based Message Authentication, Reporting and Conformance (DMARC) verification, and a summary of the actions taken for incoming messages from each domain. For more information, see the DMARC Verification Page, on page 114.
Outbreak Filtering Page	The Outbreak Filters page shows information about recent outbreaks and the messages quarantined by Outbreak Filters. Use this page to monitor your defense against phishing, scam, virus and malware attacks. For more information, see the Outbreak Filtering Page, on page 115.
URL Filtering Page	Use this page to view the URL categories most frequently occurring in messages, the most common URLs in spam messages, and the number of malicious and neutral URLs seen in messages. For more information, see the URL Filtering Page, on page 116.
Forged Email Detection Page	The Forged Email Detection report page includes the following reports: • **Top Forged Email Detection**. Displays the top ten users in the content dictionary that matched the forged From: header in the incoming messages. • **Forged Email Detection: Details**. Displays a list of all the users in the content dictionary that matched the forged From: header in the incoming messages and for a given user, the number of messages matched. For more information, see the Forged Email Detection Page, on page 118.

Email Reporting Menu	Action
File and Malware Reports	
Advanced Malware Protection Page (File Reputation and File Analysis)	The Advanced Malware Protection report page shows reporting views that displays details of Summary, File Reputation, File Analysis, File Retrospection and Mailbox Auto Remediation, for incoming and outgoing file-based threats. For more information, see the Advanced Malware Protection Page , on page 118.
Virus Filtering Page	The Virus Filtering report page provides an overview of the viruses that are sent to and from your network. This page displays the viruses that have been detected by the virus scanning engines running on the Email Security appliances and are displayed on the Security Management appliance. Use this report to take action against a particular virus. For more information, see the Virus Filtering Page, on page 123.
Macro Detection Page	The Macro Detection report page shows the top incoming and outgoing macro-enabled attachments by file type detected by the content filter and message filters. For more information, see the Macro Detection Page, on page 125.
Connection and Flow Reports	
Mail Flow Details Page	The Mail Flow Details report page provides interactive reporting on the real-time information for all remote hosts connecting to your managed Email Security appliances. You can gather information about the IP addresses, domains, and network owners (organizations) sending mail to your system. For more information, see the Mail Flow Details Page, on page 126.
Sender Groups Page	The Sender Groups report page provides a summary of connections by sender group and mail flow policy action, allowing you to review SMTP connection and mail flow policy trends. For more information, see the Sender Groups Page, on page 131.
Outgoing Destinations Page	The Outgoing Destinations report page provides information about the domains that your organization sends mail to. The top of the page includes graphs depicting the top destinations by outgoing threat messages and top destinations by outgoing clean messages. The bottom of the page displays a chart with columns sorted by total recipients (default setting). For more information, see the Outgoing Destinations Page, on page 131.
TLS Encryption Page	The TLS Encryption report page shows the overall usage of TLS connections for sent and received mail. The report also shows details for each domain sending mail using TLS connections. For more information, see the TLS Encryption Page, on page 133.

Email Reporting Menu	Action
Inbound SMTP Authentication Page	The Inbound SMTP authentication report page shows the use of client certificates and the SMTP AUTH command to authenticate SMTP sessions between the Email Security appliance and users' mail clients. For more information, see the Inbound SMTP Authentication Page, on page 136.
Rate Limits Page	The Rate Limits report page shows the mail senders (based on MAIL-FROM address) who exceed the threshold you set for the number of message recipients per sender. For more information, see the Rate Limits Page, on page 137.
Connections by Country Page	The Connections by Country report page shows the: • Top incoming mail connections based on country of origin in graphical format. • Total incoming mail connections and messages based on country of origin in tabular format. For more information, see the Connections by Country Page, on page 138.
User Reports	
User Mail Summary Page	The User Mail Summary report provides information about the mail sent and received by your internal users per email address. A single user can have multiple email addresses. The email addresses are not combined in the report. For more information, see the User Mail Summary, on page 138.
DLP Incident Summary Page	The DLP Incident Summary report page shows information on the incidents of data loss prevention (DLP) policy violations occurring in outgoing mail. For more information, see the DLP Incident Summary Page, on page 141.
Web Interaction Page	The Web Interaction report page identifies the end users who clicked URLs rewritten by policy or Outbreak Filter, and the action associated with each user click. For more information, see the Web Interaction Page, on page 142.
Filter Reports	
Message Filters Page	The Message Filters report page shows information about the top message filter matches (which message filters had the largest number of matching messages) for incoming and outgoing messages. For more information, see the Message Filters Page, on page 143.

Email Reporting Menu	Action
High Volume Mail Page	The High Volume Mail report page identifies attacks involving a large number of messages from a single sender, or with identical subjects, within a moving one-hour period. For more information, see the High Volume Mail Page, on page 144.
Content Filters Page	The Content Filters report page shows information about the top incoming and outgoing content filter matches (which content filter had the most matching messages). This page also displays the data as both bar charts and listings. For more information, see the Content Filters Page, on page 144.
System Capacity Page	The System Capacity report page shows detailed information about the overall workload of the reporting data, sent to the Security Management appliance. For more information, see the System Capacity Page, on page 145.

Mail Flow Summary Page

The Mail Flow Summary report page on the Security Management appliance provides a synopsis of the email message activity from your Email Security appliances. The Mail Flow Summary report page includes graphs and summary tables for the incoming and outgoing messages.

The Mail Flow Summary: Incoming report page shows the incoming mail graphs for the total number of messages that are processed and blocked by the appliance, as well as the summary of the incoming mails.

You can use the mail trend graphs on this page to monitor the flow of all the incoming mails that are processed and blocked by your appliances, based on the selected time range. For more information, see Choosing a Time Range for Reports , on page 26.

To search for specific information within your data, see Searching and the Interactive Email Report Pages , on page 39

The following mail trend graphs provide a visual representation of the incoming mail flow:

- Threat Detection Summary
- Content Summary

You can view the mail trend of the incoming messages based on the required counters for the respective categories. For more information, see Using Counters to Filter Data on the Trend Graphs, on page 96.

The Mail Flow Summary: Outgoing report page shows the outgoing mail graphs for the total number of messages that are processed and delivered by the appliance, as well as the summary of the outgoing mail.

You can use the mail trend graphs on this page to monitor the flow of all the outgoing mails that are processed and delivered by your appliances, based on the selected time range. For more information, see Choosing a Time Range for Reports , on page 26.

The following mail trend graphs provide a visual representation of the mail flow of the Outgoing Mails.

You can view the mail trend of the outgoing messages based on the required counters of the processed messages. For more information, see Using Counters to Filter Data on the Trend Graphs, on page 96.

The following list explains the various sections on the Mail Flow Summary report page:

Table 26: Details on the Mail Flow Summary Page

Section	Description
Mail Flow Summary: Incoming	
Number of Messages	The Number of Messages graph provides a visual representation of the total number of messages processed, including the messages that are processed as threat messages.
Threat Messages	The Threat Messages graph provides a visual representation of the total number of messages that are blocked by the Email Security appliance.
Threat Detection Summary	The Threat Detection Summary mail trend graph provides a visual representation based on the following categories: • **Connection and Reputation Filtering**: Messages that are categorized as threat by the Reputation Filtering and Invalid Recipients. • **Spam Detection**: Messages that are categorized as threat by the Anti-spam scanning engine. • **Email Spoofing**: Messages which are categorized as threat due to DMARC Verification failure. • **Outbreak Threat Summary**: Messages which are categorized as phishing, scam, virus or malware, by the Outbreak Filtering engine. • **Attachment and Malware Detection**: Messages that are categorized as threat by the Anti-virus and AMP engines. • **All Categories**: All the messages that are categorized as threat.
Content Summary	The Content Summary mail trend graph provides a visual representation based on the following categories: • **Graymail**: Messages that are categorized as marketing, bulk or social networking. • **Content Filters**: Messages that are categorized by the content filters. • **All Categories**: All the messages that are categorized by graymail engines and content filters.
Mail Flow Summary: Outgoing	

Section	Description
Number of Messages	The Number of Messages graph provides a visual representation of the total number of messages processed, including the messages that are processed as clean.
Message Delivery	The Message Delivery graph provides a visual representation of the total number of messages that are delivered, including hard bounces.
Outgoing Mails	The Outgoing Mails trend graph provides a visual representation based on the following categories: • Spam Detected • Virus Detected • Detected by AMP • Stopped by Content Filters • Stopped by DLP

How Incoming Mail Messages are Counted

Counts of incoming messages are dependent on the number of recipients per message. For example, an incoming message from example.com sent to three recipients is counted as three messages coming from that sender.

Because the messages blocked by sender reputation filtering do not actually enter the work queue, the appliance does not have access to the list of recipients for an incoming message. In this case, a multiplier is used to estimate the number of recipients. This multiplier is based on research of a large sampling of existing customer data.

How Email Messages Are Categorized by the Appliances

As messages proceed through the email pipeline, they can apply to multiple categories. For example, a message can be marked as spam or virus positive; it can also match a content filter. The precedence of the various filters and scanning activities greatly impacts the results of message processing.

In the example above, the various verdicts follow these rules of precedence:

- Spam positive
- Virus positive
- Matching a content filter

Following these rules, if a message is marked as spam positive, and your anti-spam settings are set to drop spam positive messages, the message is dropped and the spam counter is incremented.

Further, if your anti-spam settings are set to let the spam positive message continue on in the email pipeline, and a subsequent content filter drops, bounces, or quarantines the message, the spam count is still incremented. The content filter count is only incremented if the message is not spam or virus positive.

Alternately, if the message were quarantined by Outbreak Filters, it would not be counted until it was released from the quarantine and again processed through the work queue.

For complete information about message processing precedence, see chapter about the email pipeline in the online help or user guide for your Email Security appliance.

Categorizing Email Messages on the Mail Flow Summary Page

Incoming messages that are considered as threat, and outgoing messages that are delivered in the Mail Flow Summary report page are categorized as follows:

Table 27: Email Categories on Mail Flow Summary Page

Category	Description
Mail Flow Summary: Incoming	
Reputation Filtering	All connections blocked by HAT policies, multiplied by a fixed multiplier, (see the How Incoming Mail Messages are Counted, on page 46) and added with all recipients blocked by recipient throttling. The value for Stopped by Reputation Filtering is calculated based on the following factors: • Number of “throttled” messages from this sender. • Number of rejected or TCP refused connections (may be a partial count). • A conservative multiplier for the number of messages per connection. When the appliance is under heavy load, an exact count of rejected connections is not maintained on a per-sender basis. Instead, rejected connections counts are maintained only for the most significant senders in each time interval. In this situation, the value shown can be interpreted as an indicative value of the least number of messages are stopped. The Reputation Filtering total count and percentage on the Mail Flow Summary report page is always based on a complete count of all rejected connections. Only the per-sender connection counts are limited due to load.
Invalid Recipients	The total count and percentage of all mail recipients rejected by conversational LDAP rejection in addition to all RAT rejections.
Anti-Spam	The total count and percentage of incoming messages detected by the anti-spam scanning engine as positive or suspect. Additionally, messages that are both spam and virus positive.

Category	Description
Anti-Virus	The total count and percentage of incoming messages detected as virus positive and not also spam. The following messages are counted in the “Virus Detected” category: • Messages with a virus scan result of “Repaired” or “Infectious” • Messages with a virus scan result of “Encrypted” when the option to count encrypted messages as containing viruses is selected • Messages with a virus scan result of “Unscannable” when the action for unscannable messages is NOT “Deliver” • Messages with a virus scan result of “Unscannable” or “Encrypted” when the option to deliver to an alternate mail host or an alternate recipient is selected • Messages that are deleted from the Outbreak quarantine, either manually or by timing out.
Advanced Malware Protection	The total count and percentage of incoming messages blocked by the file analysis service. A message attachment was found to be malicious by file reputation filtering. This value does not include verdict updates or files found to be malicious by file analysis.
Content Filter	The total count and percentage of incoming messages that are stopped by message and content filters.
DMARC Policy	The total count and percentage of incoming messages that failed DMARC verification policy.
S/MIME Verification/Decryption Failed	The total count and percentage of incoming messages that failed S/MIME verification, decryption, or both.
Mail Flow Summary: Outgoing	
Hard Bounces	The total count and percentage of outgoing messages that are permanently undeliverable.
Delivered	The total count and percentage of outgoing messages that are delivered.

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

Additionally, if messages match a message filter and are not dropped or bounced by the filter, they are treated as clean. Messages dropped or bounced by a message filter are not counted in the totals.

Related Topics

Mail Flow Details Page, on page 126

DMARC Verification Page

The DMARC Verification report page shows the top sender domains that failed Domain-based Message Authentication, Reporting and Conformance (DMARC) verification, and a summary of the actions taken for incoming messages from each domain. You can use this report to fine-tune your DMARC settings and answer these kinds of questions:

- Which domains sent the most messages that failed DMARC verification?
- For each domain, what actions are taken on messages that failed DMARC verification?

To view the DMARC Verification report page, select DMARC Verification from the Email and Threat Reports section of the Reporting drop-down.

You can use the DMARC Verification report page to view:

- Top Domains by DMARC verification failures in graphical format.
- Total domains by DMARC verification details in tabular format. For more information, see Domains by DMARC Verification Details Table, on page 114.

If your access privileges allow you to view Message Tracking data for the messages that populate this report, click a blue number link in the table.

From the DMARC Verification report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

For more information about DMARC verification, see the Email Authentication chapter in the online help or user guide for your Email Security appliance.

Domains by DMARC Verification Details Table

The Domains by DMARC Verification Details table is an interactive table that shows the details of the sender domains that have failed (by either being rejected, quarantined, or no action), attempted, and passed the Domain-based Message Authentication, Reporting and Conformance (DMARC) verification.

To customize and sort information on the table, see Customizing Tables on Report Pages , on page 95.

To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

Outbreak Filtering Page

The Outbreak Filtering report page shows information about recent outbreaks and messages quarantined due to Outbreak Filters. You can use this page to monitor your defense against targeted virus, scam, and phishing attacks.

To view the Outbreak Filtering report page, select Outbreak Filtering from the Email Threat Reports section of the Reports drop-down.

Use the Outbreak Filtering report page to answer the following types of questions:

- How many messages are quarantined and by which Outbreak Filters rule?
- How long do messages stay in the Outbreak Quarantine?
- Which potentially malicious URLs are most frequently seen?

The following table explains the various sections on the Outbreak Filtering report page:

Table 28: Details on the Outbreak Filtering Page

Section	Description
Time Range (drop-down list)	A drop-down list with options for choosing a time range to view. For more information, see Choosing a Time Range for Reports , on page 26.
View Data For (drop-down list)	Choose an Email Security appliance for which you want to view the data, or choose All Email Appliances. See also Viewing Reporting Data for an Appliance or Reporting Group, on page 26.
Threats By Type	The Threats by Type section shows the different types of threat messages received by the appliance.
Threat Summary	The Threat Summary section shows a breakdown of the messages by Malware, Phish, Scam and Virus. To view Message Tracking details for the messages that populate this report, click a blue number link in the table.
Threat Details	The Threat Details interactive table shows details about specific outbreaks, including the threat category (virus, scam, or phishing), threat name, a description of the threat, and the number of messages identified. To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

Section	Description
Hit Messages from Incoming Messages	The Hit Messages from Incoming Messages section shows the chart and summary of the number of incoming messages processed by Outbreak Filters in the selected time period. Non-viral threats include phishing emails, scams, and malware distribution using links to an external website.
Hit Messages by Threat Level	The Hit Messages by Threat Level section shows the chart and summary of the severity of threats caught by Outbreak Filters. Level 5 threats are severe in scope or impact, while Level 1 represents low threat risk. For descriptions of threat levels, see the online help or user guide for your Email Security appliance.
Messages resided in Outbreak Quarantine	The Messages resided in Outbreak Quarantine shows the length of time messages spent in the Outbreak Quarantine. This duration is determined by the time it takes the system to compile enough data about the potential threat to make a verdict on its safety. Messages with viral threats typically spend more time in the quarantine than those with non-viral threats, because they must wait for anti-virus program updates. The maximum retention time that you specify for each mail policy is also reflected.
Top URL's Rewritten	The Top URL's Rewritten section shows the URLs that are most frequently rewritten to redirect message recipients to the Cisco Web Security Proxy for click-time evaluation of the site if and when the recipient clicks a potentially malicious link in a message. This list may include URLs that are not malicious, because if any URL in a message is deemed malicious, then all URLs in the message are rewritten. To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

Note In order to correctly populate the tables on the Outbreak Filtering report page, the appliance must be able to communicate with the Cisco update servers.

For more information, see the Outbreak Filters chapter in the online help or user guide for your Email Security appliance.

URL Filtering Page

To view the URL Filtering report page, select **URL Filtering** from the Email Threat Reports section of the Reports drop-down.

- URL Filtering report modules are populated only if URL filtering is enabled.

- URL Filtering reports are available for incoming and outgoing messages.
- Only messages that are scanned by the URL filtering engine (either as part of anti-spam/outbreak filter scanning or through message/content filters) are included in these modules. However, not all of the results are necessarily specifically attributable to the URL Filtering feature.
- The Top URL Categories module includes all categories found in messages that have been scanned, whether or not they match a content or message filter.
- Each message can be associated with only one reputation level. For messages with multiple URLs, the statistics reflect the lowest reputation of any URL in the message.
- URLs in the global whitelist configured at Security Services > URL Filtering page of the email security appliance, are not included in reports.

  URLs in whitelists used in individual filters are included in reports.
- Malicious URLs are URLs that Outbreak Filters have determined to have poor reputation. Neutral URLs are those that Outbreak Filters have determined to require click-time protection. Neutral URLs have therefore been rewritten to redirect them to the Cisco Web Security proxy.
- Results of URL category-based filters are reflected in content and message filter reports.
- Results of click-time URL evaluations by the Cisco Web Security proxy are not reflected in reports.

The following table explains the various sections on the URL Filtering report page:

Table 29: Details on the URL Filtering Page

Section	Description
Time Range (drop-down list)	A drop-down list with options for choosing a time range to view. For more information, see Choosing a Time Range for Reports , on page 26.
View Data For (drop-down list)	Choose an Email Security appliance for which you want to view the data, or choose All Email Appliances. See also Viewing Reporting Data for an Appliance or Reporting Group, on page 26.
Top URL Categories	This section displays the graphical view and summary of the top URL categories of the incoming and outgoing messages. To view Message Tracking details for the messages that populate this report, click a blue number link in the table.
Top URL Spam Messages	This section displays the graphical view and summary of the top incoming and outgoing URL spam messages.
Malicious and Neutral URLs	This section displays the chart view and the summary of malicious and neutral URLs of the incoming and outgoing messages. To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

From the URL Filtering report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

Forged Email Detection Page

To view the Forged Email Detection report page, select **Forged Email Detection** from the Email Threat Reports section of the Reports drop-down.

- The Forged Email Detection page includes the following reports:
  - **Top Forged Email Detection.** Displays the top ten users in the content dictionary that matched the forged From: header in the incoming messages.
  - **Forged Email Detection: Details.** Displays a list of all the users in the content dictionary that matched the forged From: header in the incoming messages and for a given user, the number of messages matched.
- The Forged Email Detection reports are populated only if you are using the Forged Email Detection content filter or the `forged-email-detection` message filter.

From the Forged Email Detection report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

Advanced Malware Protection Page

Advanced Malware Protection protects against zero-day and targeted file-based threats in email attachments by:

- Obtaining the reputation of known files.
- Analyzing behavior of certain files that are not yet known to the reputation service.
- Evaluating emerging threats as new information becomes available, and notifying you about files that are determined to be threats after they have entered your network.

This feature is available for incoming and outgoing messages.

For more information on the file reputation filtering and file analysis, see the *User Guide or Online Help for AsyncOS for Email Security Appliances.*

To view the Advanced Malware Protection report page, select **Advanced Malware Protection** from the Filter and Malware Reports section of the Reports drop-down.

The Advanced Malware Protection report page shows the following reporting views:

- Advanced Malware Protection – Summary, on page 119
- Advanced Malware Protection – AMP Reputation, on page 119
- Advanced Malware Protection – File Analysis, on page 120
- Advanced Malware Protection – File Retrospection, on page 120
- Advanced Malware Protection – Mailbox Auto Remediation, on page 121

See Also:

- Identifying Files by SHA-256 Hash , on page 67
- Requirements for File Analysis Report Details , on page 65
- Viewing File Reputation Filtering Data in Other Reports , on page 70

Advanced Malware Protection – Summary

The Advanced Malware Protection - Summary page shows the complete summary of the incoming and outgoing file-based threats that are identified by the file reputation and file analysis service.

For more information, see Advanced Malware Protection – AMP Reputation, on page 119 and Advanced Malware Protection – File Analysis, on page 120.

Advanced Malware Protection – AMP Reputation

The Advanced Malware Protection - AMP Reputation page shows incoming and outgoing file-based threats that were identified by the file reputation service.

For files with changed verdicts, see the AMP Verdict updates report. Those verdicts are not reflected in the Advanced Malware Protection report.

If a file extracted from a compressed or archived file is malicious, only the SHA value of the compressed or archived file is included in the Advanced Malware Protection report.

You can use the AMP Reputation view of the Advanced Malware Protection: Incoming report page to view:

- The summary of incoming files that are identified by file reputation service of the Advanced Malware Protection engine, in a graphical format.
- A trend graph for all the incoming malware threat files based on the selected time range.
- The top incoming malware threat files.
- The top incoming threat files based on the file types.
- The Incoming Malware Threat Files interactive table that lists the top incoming malware threat files.

  Drill down to view detailed analysis results, including the threat characteristics for each file.

  If your access privileges allow you to view Message Tracking data for the messages that populate this report, click a blue number link in the table.

You can use the AMP Reputation view of the Advanced Malware Protection: Outgoing report page to view:

- The summary of outgoing files that are identified by file reputation service of the Advanced Malware Protection engine, in a graphical format.
- A trend graph for all the outgoing malware threat files based on the selected time range.
- The top outgoing malware threat files.
- The top outgoing threat files based on the file types.
- The Outgoing Malware Threat Files interactive table that lists the top outgoing malware threat files that are identified by the file reputation service.

  Drill down to view detailed analysis results, including the threat characteristics for each file.

If your access privileges allow you to view Message Tracking data for the messages that populate this report, click a blue number link in the table.

Advanced Malware Protection – File Analysis

The Advanced Malware Protection - File Analysis page shows the time and verdict (or interim verdict) for each file sent for analysis. The appliance checks for analysis results every 30 minutes.

To view more than 1000 File Analysis results, export the data as a .csv file.

For deployments with an on-premises Cisco AMP Threat Grid Appliance: Files that are whitelisted on the AMP Threat Grid appliance show as "clean". For information about whitelisting, see the AMP Threat Grid documentation or online help.

Drill down to view detailed analysis results, including the threat characteristics for each file.

You can also search for additional information about an SHA, or click the link at the bottom of the file analysis details page to view additional details on the server that analyzed the file. For more information, see Identifying Files by SHA-256 Hash , on page 67.

If your access privileges allow you to view Message Tracking data for the messages that populate this report, click the **Details** link in the table.

To view details on the server that analyzed a file, see Requirements for File Analysis Report Details , on page 65.

If a file extracted from a compressed or archived file is sent for analysis, only the SHA value of the extracted file is included in the File Analysis report.

You can use the File Analysis view of the Advanced Malware Protection report page to view:

- The number of incoming and outgoing files that are uploaded for file analysis by file analysis service of the Advanced Malware Protection engine.
- A list of incoming and outgoing files that have completed file analysis requests.
- A list of incoming and outgoing files that have pending file analysis requests.

Advanced Malware Protection – File Retrospection

The Advanced Malware Protection - File Retrospection page lists the files processed by this appliance for which the verdict has changed since the message was received. For more information about on this scenario, see the documentation for your Email Security appliance.

As Advanced Malware Protection is focused on targeted and zero-day threats, threat verdicts can change as aggregated data might unviel more information.

To view more than 1000 verdict updates, export the data as a .csv file.

In the case of multiple verdict changes for a single SHA-256, this report shows only the latest verdict, not the verdict history.

To view all affected messages for a particular SHA-256 within the maximum available time range (regardless of the time range selected for the report) click a SHA-256 link.

You can use the File Retrospection view of the Advanced Malware Protection report page to view:

- A list of incoming and outgoing files with retrospective verdict changes.

Advanced Malware Protection – Mailbox Auto Remediation

The Advanced Malware Protection - Mailbox Auto Remediation report page shows the details of the mailbox remediation results for the incoming files.

You can use the Advanced Malware Protection - Mailbox Auto Remediation page to view retrospective security details such as:

- The filenames associated with a SHA-256 hash.
- Remedial actions taken on messages.
- A list of recipients for whom the mailbox remediation was successful or unsuccessful.

The Recipients for whom remediation was unsuccessful field is updated in the following scenario:

- There was a connectivity issue between your appliance and Office 365 services when the appliance was trying to perform the configured remedial action.

Click on a SHA-256 hash to view the related messages in Message Tracking.

Requirements for File Analysis Report Details

(Cloud File Analysis) Ensure That the Management Appliance Can Reach the File Analysis Server

In order to obtain File Analysis report details, the appliance must be able to connect to the File Analysis server over port 443. See details in Firewall Information, on page 401

If your Cisco Content Security Management appliance does not have a direct connection to the internet, configure a proxy server for this traffic (See Upgrade and Update Settings, on page 302.) If you have already configured the appliance to use a proxy to obtain upgrades and service updates, the existing settings are used.

If you use an HTTPS proxy, the proxy must not decrypt the traffic; use a pass-through mechanism for communications with the File Analysis server. The proxy server must trust the certificate from the Fire Analysis server, but need not provide its own certificate to the File Analysis server.

(Cloud File Analysis) Configure the Management Appliance to Display Detailed File Analysis Results

In order to allow all content security appliances in your organization to display detailed results in the cloud about files sent for analysis from any Cisco Email Security appliance or Cisco Web Security appliance in your organization, you need to join all appliances to the same appliance group.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Select **Management Appliance > Centralized Services > Security Appliances**.

Step 3 Scroll to the File Analysis section.

Step 4 If your managed appliances are pointed at different File Analysis cloud servers, select the server from which to display result details.
Result details will not be available for files processed by any other cloud server.

Step 5 Enter the Analysis Group ID.

- If you enter the Group ID incorrectly or need to change it for any other reason, you must open a case with Cisco TAC.
- This change takes effect immediately; it does not require a Commit.
- It is suggested to use your CCOID for this value.
- This value is case-sensitive.
- This value must be identical on all appliances that will share data about files that are uploaded for analysis.
- An appliance can belong to only one group.
- You can add a machine to a group at any time, but you can add it only once.

Step 6 Click **Group Now**.

Step 7 Configure the same group on each Email Security appliance that will share data with this appliance.

What to Do Next

Related Topics

For Which Files Are Detailed File Analysis Results Visible in the Cloud? , on page 70

(On-Premises File Analysis) Activate the File Analysis Account

If you have deployed an on-premises (private cloud) Cisco AMP Threat Grid Appliance, you must activate the File Analysis account for your Cisco Content Security Management appliance in order to view report details available on the Threat Grid appliance. You generally only need to do this once.

Before You Begin

Ensure that you are receiving System alerts at Critical level.

Step 1 The first time you attempt to access File Analysis report details from the Threat Grid appliance, wait a few minutes and you will receive an alert that includes a link.
If you do not receive this alert, go to **Management Appliance > System Administration > Alerts** and click **View Top Alerts**.

Step 2 Click the link in the alert message.

Step 3 Activate your management appliance account.

Additional Requirements

For any additional requirements, see the Release Notes for your Security Management appliance release, available from
http://www.cisco.com/c/en/us/support/security/content-security-management-appliance/products-release-notes-list.html

Identifying Files by SHA-256 Hash

Because filenames can easily be changed, the appliance generates an identifier for each file using a Secure Hash Algorithm (SHA-256). If an appliance processes the same file with different names, all instances are recognized as the same SHA-256. If multiple appliances process the same file, all instances of the file have the same SHA-256 identifier.

In most reports, files are listed by their SHA-256 value (in an abbreviated format).

Viewing File Reputation Filtering Data in Other Reports

Data for file reputation and analysis is available in other reports where relevant. A Detected by Advanced Malware Protection column may be hidden by default in applicable reports. To display additional columns, click the Columns link at the bottom of the table.

For Which Files Are Detailed File Analysis Results Visible in the Cloud?

If you have deployed public-cloud File Analysis, you can view detailed results for all files uploaded from any managed appliance that has been added to the appliance group for File Analysis.

If you have added your management appliance to the group, you can view the list of managed appliances in the group by clicking the button on the **Management Appliance > Centralized Services > Security Appliances** page.

Appliances in the analysis group are identified by the File Analysis Client ID. To determine this identifier for a particular appliance, look in the following location:

Appliance	Location of File Analysis Client ID
Email Security appliance	Advanced Settings for File Analysis section on the **Security Services > File Reputation and Analysis** page.
Web Security appliance	Advanced Settings for File Analysis section on the **Security Services > Anti-Malware and Reputation** page.
Cisco Content Security Management appliance	At the bottom of the **Management Appliance > Centralized Services > Security Appliances** page.

Related Topics

- (Cloud File Analysis) Configure the Management Appliance to Display Detailed File Analysis Results , on page 66

Virus Filtering Page

The Virus Filtering report page provides an overview of the viruses that are sent to and from your network. The Virus Types page displays the viruses that have been detected by the virus scanning engines running on the Email Security appliances and are displayed on the Security Management

appliance. Use this report to take action against a particular virus. For example, if you see that you are receiving a high volume of viruses known to be embedded in PDF files, you can create a filter action to quarantine messages with PDF attachments.

To view the Virus Filtering report page, select **Virus Filtering** from the File and Malware Reports section of the Reports drop-down.

If you run multiple virus scanning engines, the Virus Filtering report page includes results from all enabled virus scanning engines. The name of the virus that appears on the page is determined by the virus scanning engines. If more than one scanning engine detects a virus, it is possible to have more than one entry for the same virus.

The following list explains the various sections on the Virus Filtering report page:

Table 30: Details on the Virus Filtering Page

Section	Description
Time Range (drop-down list)	A drop-down list with options for choosing a time range to view. For more information, see Choosing a Time Range for Reports , on page 26.
View Data For (drop-down list)	Choose an Email Security appliance for which you want to view the data, or choose All Email Appliances. See also Viewing Reporting Data for an Appliance or Reporting Group, on page 26.
Top Incoming Virus Types Detected	This section displays a chart view of the detected viruses in messages sent to your network.
Top Outgoing Virus Types Detected	This section displays a chart view of the detected viruses in messages sent from your network.
Virus Types Detail	An interactive table that shows the details of each virus type. For more information, see Virus Types Detail Table, on page 125

Note To see which hosts sent virus-infected messages to your network, go to the Incoming Mail page, specify the same reporting period, and sort by virus positive messages. Similarly, to see which IP addresses have sent virus positive emails within your network, go to the Outgoing Senders page and sort by virus positive messages.

From the Virus Filtering report page, you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

You can generate a scheduled report for the Virus Filtering report page. See the Scheduling Email Reports, on page 83.

Virus Types Detail Table

The Virus Types Detail table is an interactive table that shows the total number of virus-infected messages, with a breakdown by incoming and outgoing messages. Click the column headings to sort the data.

The following table shows the table column descriptions for the Virus Types Detail table:

Table 31: Table Column Descriptions for Virus Types Detail Table

Column Name	Description
Virus Type	The name of the virus type.
Incoming Messages	Number of incoming messages detected as virus.
Outgoing Messages	Number of outgoing messages detected as virus.
Total Infected Messages	Total number of infected messages (incoming and outgoing).

Macro Detection Page

You can use the Macro Detection report page to view:

- Top Incoming Macro-Enabled Attachments by File Type in graphical and tabular format.
- Total Incoming Macro-Enabled Attachments by File Type in tabular format.
- Top Outgoing Macro-Enabled Attachments by File Type in graphical and tabular format.
- Total Outgoing Macro-Enabled Attachments by File Type in tabular format.

To view the Macro Detection report page, select **Macro Detection** from the File and Malware Reports section of the Reports drop-down.

From the Macro Detection report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

You can click on the number of macro-enabled attachments to view the related messages in Message Tracking.

Note During report generation:

- If one or more macros are detected within an archive file, the Archive Files file type is incremented by one. The number of macro-enabled attachments within an archive file are not counted.
- If one or more macros are detected within an embedded file, the parent file type is incremented by one. The number of macro-enabled attachments within an embedded file are not counted.

Mail Flow Details Page

The Mail Flow Details report page on the Security Management appliance provides interactive reporting on the real-time information for all remote hosts connecting to your managed Security Management appliances. You can gather information about the IP addresses, domains, and network owners (organizations) sending mail to your system. You can also gather information about the IP addresses and domains of the outgoing senders.

To view the Mail Flow Details report page, select **Mail Flow Details** from the Connection and Flow Reports section of the Reports drop-down.

The Mail Flow Details report page has the following tabs:

- Incoming Mail
- Outgoing Senders

To search for specific information within your data, see Searching and the Interactive Email Report Pages , on page 39.

From the Incoming Mail tab, you can:

- View the top senders by total threat messages in graphical format.
- View the top senders by clean messages in graphical format.
- View the top senders by graymail messages in graphical format.
- See the IP addresses, domains, or network owners (organizations) that have sent mail to your Security Management appliances.
- See detailed statistics on senders that have sent mail to your appliances. The statistics include the number of connections (accepted or rejected), attempted messages broken down by security service (sender reputation filtering, anti-spam, anti-virus, and so forth), total threat messages, total graymails and clean messages.
- See the Incoming Mail interactive table for the detailed information about the particular IP address, domain, or network owner (organization). For more information, see Incoming Mail Table, on page 128.

  If your access privileges allow you to view Message Tracking data for the messages that populate this report, click a number hyperlink in the table.

From the Outgoing Senders tab, you can:

- View the top senders by total threat messages in graphical format.
- View the top senders by clean messages in graphical format.
- View the top senders (by IP address or domain) of outgoing threat messages (spam, antivirus, etc.) in your organization.
- See detailed statistics on senders that have sent mail from your appliances. The statistics include the total threat messages broken down by security service (sender reputation filtering, anti-spam, anti-virus, and so forth) and clean messages.
- See the Sender Details interactive table for detailed information about the particular IP address or domain. For more information, see Sender Details Table, on page 130.

If your access privileges allow you to view Message Tracking data for the messages that populate this report, click a number hyperlink in the table.

Related Topics

- "No Domain Information" Link, on page 52
- Time Ranges in the Mail Trend Graphs, on page 53
- Views Within the Mail Flow Details Page, on page 127
- Incoming Mail Table, on page 128
- Sender Details Table, on page 130

Views Within the Mail Flow Details Page

The Mail Flow Details: Incoming report page has three different views:

- IP Addresses
- Domains
- Network Owners

These views provide a snapshot of the remote hosts connecting to the system in the context of the selected view.

Depending on the view you select, the Incoming Mail Details interactive table displays the top IP addresses, domains, or network owners that have sent mail to all public listeners configured on the Email Security appliances. You can monitor the flow of all mail into your appliances.

See the Incoming Mail Table, on page 128 for an explanation of the data included in the Incoming Mail interactive table.

From the Mail Flow Details page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

Note You can generate a scheduled report for the Mail Flow Details report page. See the Scheduling Email Reports, on page 83.

The Mail Flow Details: Outgoing report page has two different views:

- IP Addresses
- Domains

These views provide a snapshot of the remote hosts connecting to the system in the context of the selected view.

Depending on the view you select, the Sender Details interactive table displays the top IP addresses or domains of the senders that have sent mail from the public listeners configured from the Email Security appliances. You can monitor the flow of all mail from your appliances.

See the Sender Details Table, on page 130 for an explanation of the data included in the Sender Details interactive table.

"No Domain Information" Link

Domains that have connected to the Security Management appliances and could not be verified with a double-DNS lookup are automatically grouped into the special domain called "No Domain Information." You can control how these types of unverified hosts are managed via Sender Verification. For more information about Sender Verification, see the documentation or online help for your Email Security appliance.

You can use the Items Displayed menu to select the number of senders to display in the list.

Time Ranges in the Mail Trend Graphs

You can select varying degrees of granularity to see your data in a mail graph. You can select a day, week, month, and year views of the same data. Because the data is monitored in real time, information is periodically updated and summarized in the database.

For more information on time ranges, see Choosing a Time Range for Reports , on page 26.

Incoming Mail Table

The interactive Incoming Mail Details table at the bottom of the Mail Flow Details: Incoming page lists the top senders that have connected to public listeners on the Email Security appliances. The table shows domains, IP addresses, or network owners, based on the view selected.

The system acquires and verifies the validity of the remote host's IP address by performing a double DNS lookup. For more information about *double DNS lookups* and sender verification, see the user guide or online help for AsyncOS Email Security appliance.

To view Message Tracking details for the messages that populate this report, click a number hyperlink in the table.

The following table shows the table column descriptions for the Incoming Mail table:

Table 32: Table Column Descriptions for Incoming Mail Table

Column Name	Description
Sender Domain (Domains)	The domain name of the sender.
Sender IP Address (IP Addresses)	The IP address of the sender.
Hostname (IP Addresses)	The hostname of the sender.
DNS Verified (IP Addresses)	The IP addresses that are verified by the DNS.
SBRS (IP Addresses)	The SenderBase Reputation Score of the sender.
Last Sender Group (IP Addresses)	The details of the last sender group.
Last Sender Group (IP Addresses)	The details of the last sender group.

Column Name	Description
Network Owner (Network Owners)	The network owner of the sender.
Connections Rejected (Domains and Network Owners)	All connections blocked by HAT policies. When the appliance is under heavy load, an exact count of rejected connections is not maintained on a per-sender basis. Instead, rejected connections counts are maintained only for the most significant senders in each time interval.
Connections Accepted (Domains and Network Owners)	All connections accepted,
Total Attempted	All accepted and blocked connections attempted.
Stopped by Recipient Throttling (Domains and Network Owners)	This is a component of Stopped by Reputation Filtering. It represents the number of recipient messages stopped because any of the following HAT limits have been exceeded: maximum recipients per hour, maximum recipients per message, or maximum messages per connection. This is summed with an estimate of the recipient messages associated with rejected or TCP refused connections to yield Stopped by Reputation Filtering.
Stopped by Reputation Filtering	The value for Stopped by Reputation Filtering is calculated based on several factors: • Number of “throttled” messages from this sender • Number of rejected or TCP refused connections (may be a partial count) • A conservative multiplier for the number of messages per connection. When the appliance is under heavy load, an exact count of rejected connections is not maintained on a per-sender basis. Instead, rejected connections counts are maintained only for the most significant senders in each time interval. In this situation, the value shown can be interpreted as a “floor”; that is, at least this many messages are stopped. **Note** The Reputation Filtering total on the Mail Flow Summary page is always based on a complete count of all rejected connections. Only the per-sender connection counts are limited due to load.
Stopped as Invalid Recipients	All mail recipients rejected by conversational LDAP rejection plus all RAT rejections.
Spam Detected	Any spam that has been detected.
Virus Detected	Any viruses that have been detected
Detected by Advanced Malware Protection	The total count of messages detected by Advanced Malware Protection engines.
Stopped by Content Filter	The total count of messages that are stopped by a content filter.

Column Name	Description
Stopped by DMARC	The total count of messages that failed Domain-based Message Authentication, Reporting and Conformance (DMARC) verification.
Total Threat	Total number of threat messages (stopped by reputation, stopped as invalid recipient, spam, plus virus)
Marketing	Number of messages detected as unwanted marketing messages.
Social	Number of messages detected as social messages.
Bulk	Number of messages detected as bulk.
Total Graymails	Number of messages detected as graymails.
Clean	All clean messages. Messages processed on appliances on which the graymail feature is not enabled are counted as clean.

Sender Details Table

The interactive Sender Details table at the bottom of the Mail Flow Details: Outgoing page lists the top senders that have connected to public listeners on the Email Security appliances. The table shows domains or IP addresses, based on the view selected.

To view Message Tracking details for the messages that populate this report, click a number hyperlink in the table.

The following table shows the table column descriptions for the Sender Details table:

Table 33: Table Column Descriptions for Sender Details Table

Column Name	Description
Sender Domain (Domains)	The domain name of the sender.
Sender IP Address (IP Addresses)	The IP address of the sender.
Hostname (IP Addresses)	The hostname of the sender.
Spam Detected	Any spam that has been detected.
Virus Detected	Any viruses that have been detected.
Detected by Advanced Malware Protection	The total count of messages detected by Advanced Malware Protection engines.
Stopped by Content Filter	The total count of messages that are stopped by a content filter.

Column Name	Description
Stopped by DLP	The total count of messages that are stopped by DLP engine.
Total Threat	Total number of threat messages (spam, virus)
Clean	All clean messages. Messages processed on appliances on which the graymail feature is not enabled are counted as clean.
Total Messages	The total count of all the messages.

Sender Groups Page

The Sender Groups report page provides a summary of connections by sender group and mail flow policy action, allowing you to review SMTP connection and mail flow policy trends. The Mail Flow by Sender Group listing shows the percentage and number of connections for each sender group. The Connections by Mail Flow Policy Action chart shows the percentage of connections for each mail flow policy action. This page provides an overview of the effectiveness of your Host Access Table (HAT) policies. For more information about the HAT, see the documentation or online help for your Email Security appliance.

To view the Sender Groups report page, select **Sender Groups** from the Connection and Flow Reports section of the Reports drop-down.

From the Sender Group report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

Note You can generate a scheduled report for the Sender Group report page. See the Scheduling Email Reports, on page 83.

Outgoing Destinations Page

The Outgoing Destinations report page provides information about the domains that your organization sends mail to.

To view the Outgoing Destinations report page, select **Outgoing Destinations** from the Connection and Flow Reports section of the Reports drop-down.

Use the Outgoing Destinations page to answer the following types of questions:

- Which domains are the Email Security appliances sending mail to?
- How much mail is sent to each domain?
- How much of that mail is clean, spam-positive, virus-positive, malware or stopped by a content filter?
- How many messages are delivered and how many messages are hard-bounced by the destination servers?

To search for specific information within your data, see Searching and the Interactive Email Report Pages , on page 39.

The following list explains the various sections on the Outgoing Destinations report page:

Table 34: Details on the Outgoing Destinations Page

Section	Description
Time Range (drop-down list)	A drop-down list with options for choosing a time range to view. For more information, see Choosing a Time Range for Reports , on page 26.
View Data For (drop-down list)	Choose an Email Security appliance for which you want to view the data, or choose All Email Appliances. See also Viewing Reporting Data for an Appliance or Reporting Group, on page 26.
Top Destinations by Total Threat Messages	The top destination domains of outgoing threat messages (spam, antivirus, etc.) sent by your organization. Total threat messages include spam or virus positive, or the messages that are triggered by a content filter.
Top Destinations by Clean Messages	The top destination domains of clean outgoing messages sent by your organization.
Outgoing Destinations Details	All details related to the destination domains of all outgoing messages sent by your organization, sorted by total recipients. Details include detected spam, viruses, clean messages etc. For more information, see Outgoing Destinations Detail Table, on page 132. To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

From the Outgoing Destinations report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

You can generate a scheduled report for the Outgoing Destinations page. See the Scheduling Email Reports, on page 83.

Related Topics

Outgoing Destinations Detail Table, on page 132

Outgoing Destinations Detail Table

The Outgoing Destinations Detail table is an interactive table that shows the total number of messages that are processed and delivered, with a breakdown of the messages that are processed as threat (Spam, Virus, etc.) or clean, and the messages that are either hard bounced or delivered. Click the column headings to sort the data.

To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

The following table shows the table column descriptions for the Outgoing Destination Detail table:

Table 35: Table Column Descriptions for Outgoing Destination Detail Table

Column Name	Description
Destination Domain	The name of the destination domain.
Spam Detected	Number of messages detected as spam.
Virus Detected	Number of messages detected as spam.
Stopped by Content Filter	Number of messages stopped by content filters.
Total Threat	Total number of messages detected as threat (Spam, Virus, etc.)
Clean	Number of messages detected as clean,
Total Processed	Total number of messages processed as threat or clean.
Hard Bounces	Number of messages that are marked as permanently undeliverable.
Delivered	Number of messages that are delivered.
Total Messages Delivered	Total number of messages that are delivered (including Hard Bounces).

TLS Encryption Page

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

To view the TLS Encryption report page, select **TLS Encryption** from the Connection and Flow Reports section of the Reports drop-down.

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

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

The TLS Encryption report page has the following tabs:

- Incoming
- Outgoing

To search for specific information within your data, see Searching and the Interactive Email Report Pages , on page 39.

The following list explains the various sections on the TLS Encryption report page:

Table 36: Details on the TLS Encryption Page

Time Range (drop-down list)	A drop-down list with options for choosing a time range to view. For more information, see Choosing a Time Range for Reports , on page 26.
View Data For (drop-down list)	Choose an Email Security appliance for which you want to view the data, or choose All Email Appliances. See also Viewing Reporting Data for an Appliance or Reporting Group, on page 26.
TLS Connections Graph	The TLS Encryption: Incoming page displays a graph view of incoming encrypted and unencrypted TLS connections over the last hour, day, week, month or year depending on the time frame that you have selected. The TLS Encryption: Outgoing page displays a graph view of outgoing encrypted and unencrypted TLS connections over the last hour, day, week, month, or year, depending on the time frame that you have selected.
TLS Connections Summary	The TLS Encryption: Incoming page displays a table view of the total volume of incoming messages, the volume of encrypted and unencrypted messages, and the volume of successful and failed incoming TLS encrypted messages. The TLS Encryption: Outgoing page displays a table view of the total volume of outgoing messages, the volume of encrypted and unencrypted messages, and the volume of successful and failed outgoing TLS encrypted messages.
TLS Messages	The TLS Encryption: Incoming page displays a chart view of the total count and percentage of incoming TLS encrypted and unencrypted messages. The TLS Encryption: Outgoing page displays a chart view of the total count and percentage of outgoing TLS encrypted and unencrypted messages.
TLS Messages Summary	This table displays a summary of the total count and percentage of incoming and outgoing TLS encrypted and unencrypted messages.

TLS Connections Details	This table displays details for domains sending or receiving encrypted messages. For each domain, you can view the total number of connections, messages sent, and the number of TLS connections that are successful or failed. You can also view the percentage of successful and failed connections for each domain. For more information, see TLS Connections Details Table, on page 135.

TLS Connections Details Table

The TLS Connections Details table is an interactive table that shows the total number of connections, messages sent, and the number of TLS connections that are successful or failed, and the last TLS status for the incoming and outgoing messages. You can also view the percentage of successful and failed connections for each domain.

The following table shows the table column descriptions for the TLS Connection Details table:

Table 37: Table Column Descriptions for TLS Connections Details Table

Column Name	Description
Domain	The domain name of the sender.
TLS Req. Failed	All required TLS connections that failed.
TLS Req. Success	All required TLS connections that are successful.
TLS Pref. Failed	All preferred TLS connections that failed.
TLS Pref. Success	All preferred TLS connections that are successful.
Last TLS Status	The status of the TLS connections mapped based on the following: • 0: N/A • 1: Required - Fail • 2: Preferred - Fail • 3: Required - Success • 4: Preferred - Success
Total TLS Connections	Total number of TLS connections.
Unencrypted Connections	Total number of unencrypted TLS connections.
% TLS of all Connections	The percentage of TLS encryptions for all TLS connections.

Column Name	Description
Messages by TLS	The total number of TLS messages.

Inbound SMTP Authentication Page

The Inbound SMTP Authentication report page shows the use of client certificates and the SMTP AUTH command to authenticate SMTP sessions between the Email Security appliance and users' mail clients. If the appliance accepts the certificate or SMTP AUTH command, it will establish a TLS connection to the mail client, which the client will use to send a message. Since it is not possible for the appliance to track these attempts on a per-user basis, the report shows details on SMTP authentication based on the domain name and domain IP address.

To view the Inbound SMTP Authentication report page, select **Inbound SMTP Authentication** from the Connection and Flow Reports section of the Reports drop-down.

Use this report to determine the following information:

- Overall, how many incoming connection use SMTP authentication?
- How many connections use a client certificated?
- How many connections use SMTP AUTH?
- What domains are failing to connect when attempting to use SMTP authentication?
- How many connections are successfully using the fall-back when SMTP authentication fails?

The Inbound SMTP Authentication has two different views:

- Domains
- IP Addresses

These views provide a snapshot of the SMTP authentications in the context of the selected view.

The Inbound SMTP Authentication report page includes a graph for received connections, a graph for received recipients who attempted an SMTP authentication connection, and a table with details on the attempts to authenticate connections.

The following list explains the various sections on the Inbound SMTP Authentication report page:

Table 38: Details on the Inbound SMTP Authentication Page

Section	Description
Time Range (drop-down list)	A drop-down list with options for choosing a time range to view. For more information, see Choosing a Time Range for Reports , on page 26.
View Data For (drop-down list)	Choose an Email Security appliance for which you want to view the data, or choose All Email Appliances. See also Viewing Reporting Data for an Appliance or Reporting Group, on page 26.

Section	Description
Received Connection Graph	The Received Connections graph shows the incoming connections from mail clients that attempt to authentication their connections using SMTP authentication over the time range you specify. The graph displays the total number of connections the appliance received, the number that did not attempt to authenticate using SMTP authentication, the number that failed and succeeded to authenticate the connection using a client certificate, and the number that failed and succeeded to authenticate using the SMTP AUTH command.
Received Recipient Graph	The Received Recipients graph displays the number of recipients whose mail clients attempted to authenticate their connections to the Email Security appliances to send messages using SMTP authentication. The graph also show the number of recipients whose connections were authenticated and the number of recipients whose connections were not authenticated.
SMTP Authentication Details (By Domain Name or IP Address).	The SMTP Authentication Details (by domain name and IP address) table displays details about the users who attempt to authenticate their connections to the Email Security appliance to send messages. For each domain, you can view the number of connection attempts using a client certificate that were successful or failed, the number of connection attempts using the SMTP AUTH command that were successful or failed, and the number that fell back to the SMTP AUTH after their client certificate connection attempt failed.

Rate Limits Page

Rate Limiting by envelope sender allows you to limit the number of email message recipients per time interval from an individual sender, based on the mail-from address. The Rate Limits report shows you the senders who most egregiously exceed this limit.

To view the Rate Limits report page, select **Rate Limits** from the Connection and Flow Reports section of the Reports drop-down.

Use this report to help you identify the following:

- Compromised user accounts that might be used to send spam in bulk.
- Out-of-control applications in your organization that use email for notifications, alerts, automated statements, etc.
- Sources of heavy email activity in your organization, for internal billing or resource-management purposes.
- Sources of large-volume inbound email traffic that might not otherwise be considered spam.

Note that other reports that include statistics for internal senders (such as Internal Users or Outgoing Senders) measure only the number of messages sent; they do not identify senders of a few messages to a large number of recipients.

The Top Offenders by Incident chart shows the envelope senders who most frequently attempted to send messages to more recipients than the configured limit. Each attempt is one incident. This chart aggregates incident counts from all listeners.

The Top Offenders by Rejected Recipients chart shows the envelope senders who sent messages to the largest number of recipients above the configured limit. This chart aggregates recipient counts from all listeners.

Rate Limiting settings, including "Rate Limit for Envelope Senders" settings, are configured on the Email Security appliance in Mail Policies > Mail Flow Policies. For more information on rate limiting, see the documentation or online help for your Email Security appliance.

Related Topics

High Volume Mail Page, on page 144

Connections by Country Page

You can use the Connections by Country report page to view:

- Top incoming mail connections based on country of origin in graphical format.
- Total incoming mail connections and messages based on country of origin in tabular format.

To view the Connections by Country report page, select Connections by Country from the Connection and Flow Reports section of the Reports drop-down.

The following are the scenarios when no country information is displayed for the top and total incoming mail connections:

- The sender IP address belongs to a private IP address.
- The sender IP address does not get a valid SBRS.

If your access privileges allow you to view Message Tracking data for the messages that populate this report, click a blue number link in the table.

From the Connections by Country report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

User Mail Summary

The User Mail Summary report page provides information about the mail sent and received by your internal users per email address. A single user can have multiple email addresses. The email addresses are not combined in the report.

To view the User Mail Summary report page, select **User Mail Summary** from the User Reports section of the Reports drop-down.

Use the User Mail Summary interactive report page to answer these types of questions:

- Who is sending the most external email?
- Who receives the most clean email?
- Who receives the largest number of graymail messages?
- Who receives the most spam?
- Who is triggering which content filters?

- Whose email is getting caught by content filters?

To search for specific information within your data, see Searching and the Interactive Email Report Pages , on page 39.

The following list explains the various sections on the User Mail Summary report page:

Table 39: Details on the User Mail Summary Page

Section	Description
Time Range (drop-down list)	A drop-down list with options for choosing a time range to view. For more information, see Choosing a Time Range for Reports , on page 26.
View Data For (drop-down list)	Choose an Email Security appliance for which you want to view the data, or choose All Email Appliances. See also Viewing Reporting Data for an Appliance or Reporting Group, on page 26.
Top Users by Clean Incoming Messages	The top users (by domain), of clean incoming messages received by your organization.
Top Users by Clean Outgoing Messages	The top users (by domain), of clean outgoing messages sent by your organization.
Top Users by Graymail	The top users (by domain), of graymail messages.
User Mail Flow Details	The User Mail Flow Details interactive table breaks down the mails received and sent by each email address. You can sort the listing by clicking the column headers. For more information, see the User Mail Flow Details Table, on page 139. To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

From the User Mail Summary report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

Note You can generate a scheduled report for the User Mail Summary page. See the Scheduling Email Reports, on page 83.

User Mail Flow Details Table

The User Mail Flow Detail table shows detailed information about a user, including a breakdown of incoming and outgoing messages, and the number of messages in each category (such as spam detected, virus detected, stopped by content filter, etc.). Incoming and outgoing content filter matches are also shown.

Inbound Internal Users are the users for which you received email, based on the Rcpt To: address. Outbound Internal Users are based on the Mail From: address, and are useful when tracking the types of email that senders on your internal network are sending.

Some outbound mail (such as bounces) has a null sender. They are counted as outbound "unknown."

To view Message Tracking details for the messages that populate this report, click a number hyperlink in the table.

The following table shows the table column descriptions for the User Mail Flow Details table:

Table 40: Table Column Descriptions for User Mail Flow Details Table

Column Name	Description
Internal User	The domain name of the internal user.
Incoming Spam Detected	All incoming spam that was detected.
Incoming Virus Detected	The incoming virus that were detected.
Incoming Detected by Advanced Malware Protection	The incoming messages that are detected by Advanced Malware Protection (File Analysis and File Reputation).
Incoming Content Filter Matches	The incoming content filter matches that were detected.
Incoming Stopped by Content Filter	The incoming messages that were stopped due to content filters that have been set.
Incoming Marketing	The incoming messages that were detected as marketing.
Incoming Social Networking	The incoming messages that were detected as social networking.
Incoming Bulk	The incoming messages that were detected as bulk.
Incoming Graymails	The incoming messages that were detected as graymail.
Incoming Clean	All incoming clean messages.
Outgoing Spam Detected	The outgoing spam that was detected.
Outgoing Virus Detected	The outgoing viruses that were detected.
Outgoing Content Filter Matches	The outgoing content filter matches that were detected.
Outgoing Stopped by Content Filter	The outgoing messages that were stopped due to content filters that have been set.
Outgoing Clean	All outgoing clean messages.

Searching for a Specific Internal User

With the search form at the bottom of the User Mail Summary page and the User Mail Flow Details page, you can search for a specific internal user (email address). Select whether to exactly match the search text or look for items starting with the entered text (for example, starts with “ex” will match “example@example.com”).

DLP Incident Summary Page

The DLP Incidents (DLP Incident Summary) report page shows information on the incidents of data loss prevention (DLP) policy violations occurring in outgoing mail. The Email Security appliance uses the DLP email policies enabled in the Outgoing Mail Policies table to detect sensitive data sent by your users. Every occurrence of an outgoing message violating a DLP policy is reported as an incident.

To view the DLP Incident Summary report page, select **DLP Incidents** from the User Reports section of the Reports drop-down.

Using the DLP Incident Summary report, you can answer these kinds of questions:

- What type of sensitive data is being sent by your users?
- How severe are these DLP incidents?
- How many of these messages are being delivered?
- How many of these messages are being dropped?
- Who is sending these messages?

The DLP Incident Summary page contains two main sections:

- The DLP incident trend graphs summarizing the top DLP incidents by severity (Low, Medium, High, Critical) and policy matches.
- The DLP Incident Details listing.

From the DLP Incidents report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

To search for specific information within your data, see Searching and the Interactive Email Report Pages , on page 39.

The following list explains the various sections on the DLP Incident Summary report page:

Table 41: Details on the DLP Incident Summary Page

Section	Description
Time Range (drop-down list)	A drop-down list with options for choosing a time range to view. For more information, see Choosing a Time Range for Reports , on page 26.
View Data For (drop-down list)	Choose an Email Security appliance for which you want to view the data, or choose All Email Appliances. See also Viewing Reporting Data for an Appliance or Reporting Group, on page 26.

Section	Description
Top Incidents by Severity	The top DLP incidents listed by severity.
Incident Summary	The DLP policies currently enabled for each email appliance's outgoing mail policies are listed in the DLP Incident Details interactive table at the bottom of the DLP Incident Summary page. Click the name of a DLP policy to view more detailed information.
Top DLP Policy Matches	The top DLP Policies that have been matched.
DLP Incident Details	The DLP Incidents Details table shows the total number of DLP incidents per policy, with a breakdown by severity level, and whether any of the messages are delivered in the clear, delivered encrypted, or dropped. To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

Web Interaction Page

To view the Web Interaction report page, select **Web Interaction** from the User Reports section of the Reports drop-down.

- Web Interaction report modules are populated only if the Web Interaction Tracking feature is enabled on managed Email Security appliances.
- Web Interaction reports are available for incoming and outgoing messages.
- Only rewritten URLs (either by policy or Outbreak Filter) clicked by the end users are included in these modules.
- Web Interaction Tracking page includes the following reports:
  - Top Malicious URLs clicked by End Users.
  - Top Users who clicked on Rewritten Malicious URLs.
  - Web Interaction Tracking Details.

From the Web Interaction report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

The following list explains the various sections on the Web Interaction report page:

Table 42: Details on the Web Interaction Page

Section	Description
Time Range (drop-down list)	A drop-down list with options for choosing a time range to view. For more information, see Choosing a Time Range for Reports , on page 26.

Section	Description
View Data For (drop-down list)	Choose an Email Security appliance for which you want to view the data, or choose All Email Appliances. See also Viewing Reporting Data for an Appliance or Reporting Group, on page 26.
Top Malicious URLs clicked by End Users	This section displays the summary of the top malicious URLs clicked by end users, for incoming and outgoing messages.
Top Users who clicked on Malicious URLs	This section displays the summary of the top end users who clicked on the Rewritten Malicious URLs, for incoming and outgoing messages.
Web Interaction Tracking Details	This section displays the chart view and the summary of malicious and neutral URLs of the incoming and outgoing messages. To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

Web Interaction Tracking Details

The Web Interaction Tracking Details table is an interactive table which includes the following information:

- A list of all the rewritten URLs (malicious and unmalicious).
- Action taken (allow, block, or unknown) when a rewritten URL was clicked.
- If the verdict of a URL (clean or malicious) was unknown at the time when the end user clicked it, the status is shown as unknown. This could be because the URL was under further scrutiny or the web server was down or not reachable at the time of the user click.
- The number of times end users clicked on a rewritten URL.
- Note the following:
  - If you have configured a content or message filter to deliver messages after rewriting malicious URLs and notify another user (for example, an administrator), the web interaction tracking data for the original recipient is incremented if the notified user clicks on the rewritten URLs.
  - If you are sending a copy of quarantined messages containing rewritten URLs to a user other than the original recipient (for example, to an administrator) using the web interface, the web interaction tracking data for the original recipient is incremented if the other user clicks on the rewritten URLs.

To view Message Tracking details for the messages that populate this report, click a blue number link in the table.

Message Filters Page

The Message Filters report page shows information about the top message filter matches (which message filters had the largest number of matching messages) for incoming and outgoing messages.

You can use the Message Filters report page to view:

- Top message filter by number of matches in graphical format.
- Total message filter by number of matches in tabular format.

To view the Message Filters report page, select **Message Filters** from the Filter Reports section of the Reports drop-down.

From the Message Filters report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

High Volume Mail Page

Use reports on this page to:

- Identify attacks involving a large number of messages from a single sender, or with identical subjects, within a moving one-hour period.
- Monitor top domains to ensure that such attacks do not originate in your own domain. If this situation occurs, one or more accounts in your organization may be compromised.
- Help identify false positives so you can adjust your filters accordingly.

To view the High Volume Mail report page, select **High Volume Mail** from the Filter Reports section of the Reports drop-down.

You can use the High Volume Mail report page to view:

- Messages with the top subjects in graphical format.
- Messages with the top envelope senders in graphical format.
- Top message filters by number of matches in graphical format.
- Total message filters by number of matches in tabular format.

From the High Volume Mail report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

Reports on this page show data only from message filters that use the Header Repeats rule and that pass the number-of-messages threshold that you set in that rule. When combined with other rules, the Header Repeats rule is evaluated last, and is not evaluated at all if the message disposition is determined by a preceding condition. Similarly, messages caught by Rate Limiting never reach Header Repeats message filters. Therefore, some messages that might otherwise be considered high-volume mail may not be included in these reports. If you have configured your filters to whitelist certain messages, those messages are also excluded from these reports.

For more information about message filters and the Header Repeats rule, see the online help or user guide for your Email Security appliance.

Content Filters Page

The Content Filters report page shows information about the top incoming and outgoing content filter matches (which content filter had the most matching messages). The page displays the data as both bar charts and listings. Using the Content Filters report page, you can answer the following types of questions:

- Which content filter is triggered the most by incoming or outgoing mail?
- Who are the top users sending or receiving mail that triggers a particular content filter?

To view the Content Filters report page, select Content Filters from the Filter Reports section of the Reports drop-down.

You can use the Content Filter report page to view:

- Top incoming and outgoing content filter matches in graphical format.
- Top incoming and outgoing content filter matches in tabular format.

From the Content Filters report page you can export raw data to a CSV file. For information on printing or exporting a file, see the Printing and Exporting Reporting Data , on page 31.

Note Note You can generate a scheduled report for the Content Filter page. See the Scheduling Email Reports, on page 83.

System Capacity Page

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

The System Capacity report page can be used to determine the following information:

- Identify when Email Security appliances are exceeding recommended capacity; this enables you to determine when configuration optimization or additional appliances are needed.
- Identify historical trends in system behavior that point to upcoming capacity issues.
- For troubleshooting, identify which parts of the system are using the most resources.

Monitor your Email Security appliances to ensure that the capacity is appropriate to your message volumes. Over time, volume inevitably rises and appropriate monitoring ensures that additional capacity or configuration changes can be applied proactively. The most effective way to monitor system capacity is to track the overall volume, the messages in the work queue, and the incidents of Resource Conservation Mode.

- **Volume**: It is important to understand the “normal” message volume and the “usual” spikes in your environment. Track this data over time to measure volume growth. You can use the Incoming Mail and Outgoing Mail pages to track volume over time. For more information, see System Capacity – Incoming Mail, on page 77 and System Capacity – Outgoing Mail, on page 78.
- **Work Queue**: The work queue is designed to work as a “shock absorber”— absorbing and filtering spam attacks and processing unusual increases in non-spam messages. However, the work queue can also indicate a system under stress. Prolonged and frequent work queue backups may indicate a capacity problem. You can use the System Capacity – Workqueue page to track the activity in your work queue. For more information, see System Capacity – Workqueue, on page 77.
- **Resource Conservation Mode**: When an appliance becomes overloaded, it enters Resource Conservation Mode (RCM) and sends a CRITICAL system alert. This is designed to protect the device and allow it to process any backlog of messages. Your appliance should enter RCM infrequently and only during a

very large or unusual increase in mail volume. Frequent RCM alerts may be an indication that the system is becoming overloaded. See Resource Conservation Activity , on page 78.

How to Interpret the Data You See on System Capacity Page

When choosing time ranges for viewing data on the System Capacity page, the following is important to remember:

- Day Report— The Day report queries the hour table and displays the exact number of queries that have been received by the appliance on an hourly basis over a 24 hour period. This information is gathered from the hour table. This is an exact number.
- Month Report— The Month report queries the day tables for the 30 or 31 days (dependent on the number of days in the month), giving you an exact report on the number of queries over 30 or 31 days. Again, this is an exact number.

The 'Maximum' value indicator on the System Capacity page is the highest value seen for the specified period. The 'Average' value is the average of all values for the specified period. The period of aggregation depends on the interval selected for that report. For example, you can choose to see the Average and Maximum values for each day if the chart is for a month period.

You can click the View Details link for a specific graph to view data for individual Email Security appliances and overall data for the appliances connected to the Security Management appliance.

System Capacity – Workqueue

The Workqueue page shows the average time a message spends in the work queue, excluding any time spent in the Spam quarantine or in a policy, virus, or outbreak quarantine. You can view time periods from an hour up to one month. This average can help in identifying both short term events delaying mail delivery and identify long term trends in the workload on the system.

Note If a message is released from the quarantine into the work queue, the "average time in work queue" metric ignores this time. This prevents double-counting and distorted statistics due to extended time spent in a quarantine.

The report also shows the volume of messages in the work queue over a specified time period, and it shows the maximum messages in the work queue over the same time period. The graphical representation of the maximum messages in the work queue also shows the work queue threshold level.

Occasional spikes in the Workqueue graphs are normal and expected. If the messages in the work queue remain higher than the configured threshold for a long duration, this may indicate a capacity issue. In this scenario, consider tuning the threshold level or review the system configuration.

To change the work queue threshold level, see Adjusting the Reference Threshold in System Health Graphs for Email Security Appliances, on page 337.

Tip When reviewing the work queue page, you may want to measure the frequency of work queue backups, and take note of work queue backups that exceed 10,000 messages.

System Capacity – Incoming Mail

The System Capacity – Incoming Mail page shows incoming connections, the total number of incoming messages, the average message size, and the total incoming message size. You can view the results for a day, week, month, or year. It is important to understand the trends of normal message volume and spikes in your environment. You can use the System Capacity – Incoming Mail page to track volume growth over time and plan for system capacity. You might also want to compare the incoming mail data with the sender profile data to view the trends in volumes of email messages that are sent from specific domains to your network.

Note An increased number of incoming connections may not necessarily affect system load.

System Capacity – Outgoing Mail

The System Capacity – Outgoing Mail page shows outgoing connections, the total number of outgoing messages, the average message size, and the total outgoing message size. You can view the results for a day, week, month, or year. It is important to understand the trends of normal message volume and spikes in your environment. You can use the System Capacity – Outgoing Mail page to track volume growth over time and plan for system capacity. You might also want to compare the outgoing mail data with the outgoing destinations data to view the trends in volumes of email messages that are sent from specific domains or IP addresses.

System Capacity – System Load

The system load report shows the following:

Overall CPU Usage

Email Security appliances are optimized to use idle CPU resources to improve message throughput. High CPU usage may not indicate a system capacity problem. If the high CPU usage is coupled with consistent, high-volume memory page swapping, you may have a capacity problem.

Note This graph also indicates a threshold for CPU usage that is a visual reference only. To adjust the position of this line, see Adjusting the Reference Threshold in System Health Graphs for Email Security Appliances, on page 337. You can configure your Email Security appliances to send you alerts that will suggest actions that you can take to address capacity issues.

This page also shows a graph that displays the amount of CPU used by different functions, including mail processing, spam and virus engines, reporting, and quarantines. The CPU-by-function graph is a good indicator of which areas of the product use the most resources on your system. If you need to optimize your appliance, this graph can help you determine which functions may need to be tuned or disabled.

Memory Page Swapping

The memory page swapping graph shows how frequently the system must page to disk, in kilobytes per second.

The system is designed to swap memory regularly, so some memory swapping is expected and is not an indication of problems with your appliance. Unless the system *consistently* swaps memory in high volumes, memory swapping is normal and expected behavior (especially on C170appliances). To improve performance,

you may need to add Email Security appliances to your network or tune your configuration to ensure maximum throughput.

Note This graph also indicates a threshold for memory page swapping that is a visual reference only. To adjust the position of this line, see Adjusting the Reference Threshold in System Health Graphs for Email Security Appliances, on page 337. You can configure your Email Security appliances to send you alerts that will suggest actions that you can take to address capacity issues.

Resource Conservation Activity

The resource conservation activity graph shows the number of times the Email Security appliance entered Resource Conservation Mode (RCM). For example, if the graph shows n times, it means that the appliance has entered RCM n times and exited at least n-1 times.

Your appliances should enter RCM infrequently and only during a very large or unusual increase in mail volume. If the Resource Conservation Activity graph shows that your appliance is entering RCS frequently, it may be an indication that the system is becoming overloaded.

System Capacity – All

The **All** page consolidates all the previous system capacity reports onto a single page so you can view the relationship between the different reports. For example, you might see that the message queue is high at the same time that excessive memory swapping takes place. This might be an indication that you have a capacity problem. You may want to save this page as a PDF file to preserve a snapshot of system performance for later reference (or to share with support staff).

Threshold Indicator in System Capacity Graphs

In some graphs, a line indicates the default value that may indicate a possible problem if it is frequently or consistently crossed. To adjust this visual indicator, see Adjusting the Reference Threshold in System Health Graphs for Email Security Appliances, on page 337.

Reporting Data Availability Page

The **Email > Reporting > Reporting Data Availability** page allows you to view, update and sort data to provide real-time visibility into resource utilization and email traffic trouble spots.

All data resource utilization and email traffic trouble spots are shown from this page, including the data availability for the overall appliances that are managed by the Security Management appliance.

From this report page you can also view the data availability for a specific appliance and time range.

Reporting of Graymail

Graymail statistics are reflected in the following reports:

Report	Contains the Following Graymail Data
Mail Flow Summary page > Incoming tab	The number of incoming graymail messages under each graymail category (Marketing, Social, and Bulk) and the total number of graymail messages.
Mail Flow Details page > Outgoing Senders tab	The top graymail senders.
Mail Flow Details page > Incoming Mails tab	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.
User Mail Summary page > Top Users by Graymail	The top end users who receive graymail.
User Mail Summary page > User 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 users.

Reporting of Marketing Messages after Upgrade to AsyncOS 9.5

After upgrade to AsyncOS 9.5:

- The number of marketing messages is the 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.
- If the graymail feature is not enabled on managed Email Security appliances, marketing messages are counted as clean messages.

About Scheduled and On-Demand Email Reports

Types of Reports Available

Except as noted, the following types of Email Security reports are available as both scheduled and on-demand reports:

- Content Filters—This report includes up to 40 content filters. For additional information on what is included on this page, see the Content Filters Page, on page 144.
- DLP Incident Summary—For information on what is included on this page, see the DLP Incident Summary Page, on page 141.

- Delivery Status—The report page displays information about delivery problems to a specific recipient domain or Virtual Gateway address, page displays a list of the top 20, 50, or 100 recipient domains for messages delivered by the system within the last three hours. You can sort by latest host status, active recipients (the default), connections out, delivered recipients, soft bounced events, and hard bounced recipients by clicking the links in the column heading for each statistic. For more information on what the Delivery Status page does on the Email Security appliance, see the documentation or online help for your Email Security appliance.
- Domain-Based Executive Summary—This report is based on the Mail Flow Summary Page, on page 109, and is limited to a group of specified domains. For information on what is included, see the Domain-Based Executive Summary Report, on page 81.
- Executive Summary—This report is based on the information from the Mail Flow Summary Page, on page 109. For information on what is included, see the Domain-Based Executive Summary Report, on page 81.
- Mail Flow Details — For information on what in included on this page, see the Mail Flow Details Page, on page 126.
- User Mail Summary—For information on what is included on this page, see the User Mail Summary, on page 138.
- Outgoing Destinations—For information on what is included on this page, see the Outgoing Destinations Page, on page 131.
- Sender Groups —For information on what is included on this page, see the Sender Groups Page, on page 131.
- TLS Encrptions—For information on what is included on this page, see the TLS Encryption Page, on page 133.
- Virus Types—For information on what is included on this page, see the Virus Filtering Page, on page 123.

Time ranges

Depending on the report, these reports can be configured to include data for the previous day, previous seven days, previous month, previous calendar day (up to 250), or previous calendar month (up to 12). Alternatively, you can include data for a custom number of days (from 2 days to 100 days) or a custom number of months (from 2 months to 12 months).

Regardless of when you run a report, the data is returned from the previous time interval (hour, day, week, or month). For example, if you schedule a daily report to run at 1AM, the report will contain data from the previous day, midnight to midnight (00:00 to 23:59).

Languages and Locales

Note You can schedule a PDF report or export raw data as a CSV file with a specific locale for that individual report. The language drop-down menu on the Scheduled Reports page allows you to view or schedule a PDF report in the users current selected locale and language. See important information at Printing and Exporting Reporting Data , on page 31.

Storage of Archived Reports

For information on how long reports are stored for, and when archived reports are deleted from the system, see Viewing and Managing Archived Email Reports , on page 87.

Additional Report Types

Two special reports that can be generated in the **Email > Reporting** section on the Security Management appliance are:

Domain-Based Executive Summary Report

The Domain-Based Executive Summary report provides a synopsis of the incoming and outgoing message activity for one or more domains in your network. It is similar to the Executive Summary report, but it limits the report data to the messages sent to and from the domains that you specify. The outgoing mail summary shows data only when the domain in the PTR (pointer record) of the sending server matches a domain you specify. If multiple domains are specified, the appliance aggregates the data for all those domains into a single report.

To generate reports for a subdomain, you must add its parent domain as a second-level domain in the reporting system of the Email Security appliance and the Security Management appliance. For example, if you add example.com as a second-level domain, its subdomains, such as subdomain.example.com, are available for reporting. To add second-level domains, use **reportingconfig -> mailsetup -> tld** in the Email Security appliance CLI, and **reportingconfig -> domain -> tld** in the Security Management appliance CLI.

Unlike other scheduled reports, Domain-Based Executive Summary reports are not archived.

Domain-Based Executive Summary Reports and Messages Blocked by Sender Reputation Filtering

Because messages blocked by sender reputation filtering do not enter the work queue, AsyncOS does not process these messages to determine the domain destination. An algorithm estimates the number of rejected messages per domain. To determine the exact number of blocked messages per domain, you can delay HAT rejections on the Security Management appliance until the messages reach the recipient level (RCPT TO). This allows AsyncOS to collect recipient data from the incoming messages. You can delay rejections using **listenerconfig -> setup** command on the Email Security appliance. However, this option can impact system performance. For more information about delayed HAT rejections, see the documentation for your Email Security appliance.

Note To see Stopped by Reputation Filtering results in your Domain-Based Executive Summary report on the Security Management appliance, you must have **hat_reject_info** enabled on both the Email Security appliance and the Security Management appliance.To enable the **hat_reject_info** on the Security Management appliance, run the **reportingconfig > domain > hat_reject_info** command.

Managing Lists of Domains and Recipients for Domain-Based Executive Summary Reports

You can use a configuration file to manage the domains and recipients for a Domain-Based Executive Summary report. The configuration file is a text file that is stored in the configuration directory of the appliance. Each line in the file produces a separate report. This allows you to include a large number of domains and recipients in a single report, as well as define multiple domain reports in a single configuration file.

Each line of the configuration file includes a space-separated list of domain names and a space-separated list of email addresses for the report recipients. A comma separates the list of domain names from the list of email addresses. You can include subdomains by appending the subdomain name and a period at the beginning of the parent domain name, such as subdomain.example.com.

The following is a Single Report configuration file that generates three reports.

```
yourdomain.com sampledomain.com, admin@yourdomain.com
sampledomain.com, admin@yourdomain.com user@sampledomain.com
subdomain.example.com mail.example.com, user@example.com
```

Note You can use a configuration file and the settings defined for a single named report to generate multiple reports at the same time. For example, a company named Bigfish purchases two other companies, Redfish and Bluefish, and continues to maintain their domains. Bigfish creates a single Domain-Based Executive Summary report using a configuration file containing three lines corresponding to separate domain reports. When the appliance generates a Domain-Based Executive Summary report, an administrator for Bigfish receives a report on the Bigfish.com, Redfish.com, and Bluefish.com domains, while a Redfish administrator receives a report on the Redfish.com domain and a Bluefish administrator receives a report on the Bluefish.com domain.

You can upload a different configuration file to the appliance for each named report. You can also use the same configuration file for multiple reports. For example, you might create separate named reports that provide data about the same domains over different time periods. If you update a configuration file on your appliance, you do not have to update the report settings in the GUI unless you change the filename.

Creating Domain-Based Executive Summary Reports

Step 1 On the Security Management appliance, you can schedule the report or generate the report immediately.
To schedule the report:

a) [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.
b) Choose **Email > Reporting > Scheduled Reports**.
c) Click **Add Scheduled Report**.
To create an on-demand report:

- [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.
- Choose **Email > Reporting > Archived Reports.**
- Click **Generate Report Now**.

Step 2 From the **Report Type** drop-down list, choose **Domain-Based Executive Summary** report type.

Step 3 Specify the domains to include in the report and the email addresses for the report recipients. You can select one of the following options for generating the report:

- Generate report by specifying individual domains. Enter the domains for the report and the email addresses for the report recipients. Use commas to separate multiple entries. You can also use subdomains, such as

subdomain.yourdomain.com. Specifying individual domains is recommended if you create reports for a small number of domains that are not expected to change frequently.

- Generate reports by uploading file. Import a configuration file that contains a list of the domains and recipient email addresses for the report. You can select a configuration file from the configuration directory on the appliance or upload one from your local computer. Using a configuration file is recommended if you create reports for a large number of domains that change frequently. For more information on configuration files for domain-based reports, see Managing Lists of Domains and Recipients for Domain-Based Executive Summary Reports , on page 81.

Note If you send reports to an external account (such as Yahoo! Mail or Gmail), you may need to add the reporting return address to the external account's whitelist to prevent report messages from being incorrectly classified as spam.

Step 4 In the Title text field, type the name of the title for the report.
AsyncOS does not verify the uniqueness of report names. To avoid confusion, do not create multiple reports with the same name.

Step 5 In the Outgoing Domain section, choose the domain type for the outgoing mail summary. Choices are: By Server or By Email Address.

Step 6 From the Time Range to Include drop-down list, select a time range for the report data.

Step 7 In the Format section, choose the format of the report.
Choices include:

- PDF. Create a formatted PDF document for delivery, archival, or both. You can view the report as a PDF file immediately by clicking Preview PDF Report.
- CSV. Create an ASCII text file that contains raw data as comma-separated values. Each CSV file may contain up to 100 rows. If a report contains more than one type of table, a separate CSV file is created for each table.

Step 8 From the Schedule section, choose a schedule for generating the report.
Choices include: Daily, Weekly (drop-down list for day of week included), or monthly.

Step 9 (Optional) Upload a custom logo for the report. The logo appears at the top of the report.

- The logo should be a .jpg, .gif, or .png file that is at most 550 x 50 pixels.
- If a logo file is not supplied, the default Cisco logo is used.

Step 10 Select a language for this report. For generating PDFs in Asian languages, see important information at Printing and Exporting Reporting Data , on page 31.

Step 11 Click **Submit** to submit your changes on the page, then click **Commit Changes** to commit your changes.

Executive Summary Report

The Executive Summary Report is a high-level overview of the incoming and outgoing email message activity from your Email Security appliances. that can be viewed on the Security Management appliance.

This report page summarizes what you can view on the Mail Flow Summary Page, on page 109. For more information on the Email Reporting Overview page, see Mail Flow Summary Page, on page 109.

Scheduled Reports Page

• Scheduling Email Reports, on page 83

Scheduling Email Reports

You can schedule any of the reports listed in About Scheduled and On-Demand Email Reports , on page 79.

To manage report scheduling, see the following:

Adding Scheduled Reports

To add a scheduled email report, use the following steps:

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Scheduled Reports**.

Step 3 Click **Add Scheduled Report**.

Step 4 Choose your report type.
For descriptions of the report types, see About Scheduled and On-Demand Email Reports , on page 79.

Note - For information about the settings for a Domain-Based Executive Summary report, see Domain-Based Executive Summary Report, on page 81.

- Available options for scheduled reports differ by report type. Options described in the remainder of this procedure do not necessarily apply to all reports.

Step 5 In the **Title** field, type the title of your report.
To avoid creating multiple reports with the same name, we recommend using a descriptive title.

Step 6 Choose the time range for the report from the **Time Range to Include** drop-down menu.

Step 7 Choose the format for the generated report.
The default format is PDF. Most reports also allow you to save raw data as a CSV file.

Step 8 Depending on the report, for Number of Rows, choose the amount of data to include.

Step 9 Depending on the report, choose the column by which to sort the report.

Step 10 From the **Schedule** area, select the radio button next to the day, week, or month for your scheduled report. Additionally, include the time that you want the report scheduled for. Time increments are based on midnight to midnight (00:00 to 23:59).

Step 11 In the **Email** text field, type in the email address where the generated report will be sent.
If you do not specify an email recipient, the system will still archive the reports.

You can add as many recipients for reports as you want, including zero recipients. If you need to send the reports to a large number of addresses, however, you may want to create a mailing list instead of listing the recipients individually.

Step 12 Choose a language for the report.
For Asian languages, see important information at Printing and Exporting Reporting Data , on page 31.

Step 13 Click **Submit**.

Editing Scheduled Reports

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Scheduled Reports**.

Step 3 Click the report name link in the Report Title column that you want to modify.

Step 4 Modify the report settings.

Step 5 Submit and commit your changes.

Discontinuing Scheduled Reports

To prevent future instances of scheduled reports from being generated, perform the following steps:

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Scheduled Reports**.

Step 3 Select the check boxes corresponding to the reports that you want to discontinue generating. To remove all scheduled reports, select the **All** check box.

Step 4 Click **Delete**.

Note Any archived versions of deleted reports are *not* automatically deleted. To delete previously-generated reports, see Deleting Archived Reports, on page 87.

Generating Email Reports On Demand

In addition to the reports that you can view (and generate PDFs for) using the interactive report pages described in Understanding the Email Reporting Pages, on page 106, you can save PDFs or raw-data CSV files for the reports listed in About Scheduled and On-Demand Email Reports , on page 79 at any time, for the time frame that you specify.

To generate an on-demand report perform the following:

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Archived Reports**.

Step 3 Click **Generate Report Now**.

Step 4 Choose a report type.
For descriptions of the report types, see About Scheduled and On-Demand Email Reports , on page 79.

Step 5 In the Title text field, type the name of the title for the report.
AsyncOS does not verify the uniqueness of report names. To avoid confusion, do not create multiple reports with the same name.

Note For information about the settings for a Domain-Based Executive Summary report, see Domain-Based Executive Summary Report, on page 81.

Available options for scheduled reports differ by report type. Options described in the remainder of this procedure do not necessarily apply to all reports.

Step 6 From the Time Range to Include drop-down list, select a time range for the report data.
Note the custom time range option.

Step 7 In the Format section, choose the format of the report.
Choices include:

- PDF. Create a formatted PDF document for delivery, archival, or both. You can view the report as a PDF file immediately by clicking Preview PDF Report.
- CSV. Create an ASCII text file that contains raw data as comma-separated values. Each CSV file may contain up to 100 rows. If a report contains more than one type of table, a separate CSV file is created for each table.

Step 8 Select the appliances or appliance groups for which you want to run the report. If you have not created any appliance groups, this option does not appear.

Step 9 From the Delivery Option section, choose the following:

- Archive the report by checking the **Archive Report** checkbox.

By choosing this, the report will be listed on the Archived Reports page.

Note Domain-Based Executive Summary reports cannot be archived.

- Email the report, by checking the **Email now to recipients** checkbox.

In the text field, type in the recipient email addresses for the report.

Step 10 Select a language for this report. For generating PDFs in Asian languages, see important information at Printing and Exporting Reporting Data , on page 31.

Step 11 Click **Deliver This Report** to generate the report.

Archived Email Reports Page

Viewing and Managing Archived Email Reports

Scheduled and on-demand reports are archived for a period of time.

The Security Management appliance retains the most recent reports that it generates, up to 30 instances of each scheduled report, up to 1000 total versions for all reports. The limit of 30 instances applies to each scheduled report with the same name and time range.

Archived reports are deleted automatically. As new reports are added, older reports are removed to keep the number at 1000.

Archived reports are stored in the /periodic_reports directory on the appliance. (See IP Interfaces and Accessing the Appliance, on page 389 for more information.)

Accessing Archived Reports

The **Email > Reporting > Archived Reports** page lists scheduled and on-demand reports that you have chosen to archive which have been generated and not yet purged.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Archived Reports**.

Step 3 To locate a particular report if the list is long, filter the list by choosing a report type from the **Show** menu, or click a column heading to sort by that column.

Step 4 Click the Report Title to view that report.

Deleting Archived Reports

Reports are automatically deleted from the system according to the rules outlined in Viewing and Managing Archived Email Reports , on page 87. However, you can manually delete unneeded reports.

To manually delete Archived reports, perform the following:

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Reporting > Archived Reports**.
The Archived reports that are available are displayed.

Step 3 Select the checkbox for one or more reports to delete.

Step 4 Click **Delete**.

Step 5 To prevent future instances of scheduled reports from being generated, see Discontinuing Scheduled Reports , on page 85.

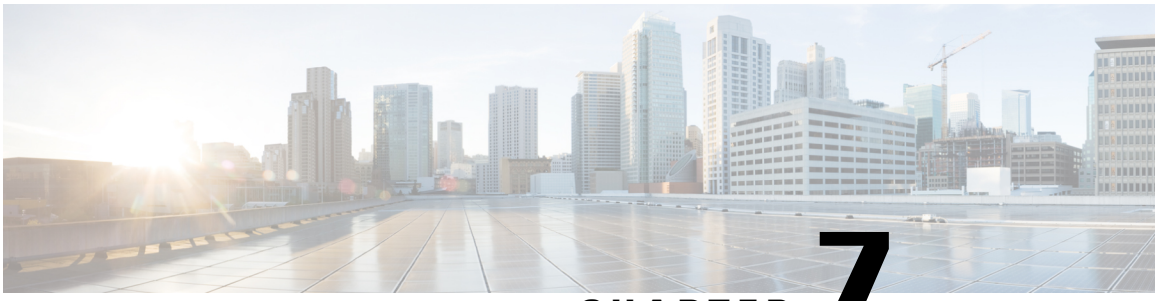

CHAPTER 7

Tracking Messages

This chapter contains the following sections:

Tracking Service Overview

The tracking service of the Cisco Content Security Management appliance complements Email Security appliances. With the Security Management appliance, email administrators have a single place to track the status of messages that traverse any of their Email Security appliances.

The Security Management appliance makes it easy to find the status of messages that Email Security appliances process. Email administrators can quickly resolve help desk calls by determining the exact location of a message. With the Security Management appliance, an administrator can determine if a particular message was delivered, found to contain a virus, or placed in a spam quarantine — or if it is located somewhere else in the mail stream.

Instead of having to search through log files using grep or similar tools, you can use the flexible tracking interface of the Security Management appliance to locate messages. You can use a variety of search parameters in combination

Tracking queries can include:

- **Time Frame:** Find a message that was sent between specified dates and times.
- **Envelope Information:** Find messages from particular envelope senders or recipients by entering the text strings to match.
- **Subject:** Match a text string in the subject line. Warning: Do not use this type of search in environments where regulations prohibit such tracking.

- **Attachment Name**: You can search for messages based on an attachment name. Messages that contain at least one attachment with the queried name will appear in the search results.

  For performance reasons, the names of files within attachments such as OLE objects or archives such as .ZIP files are not tracked.

  Some attachments may not be tracked. For performance reasons, scanning of attachment names occurs only as part of other scanning operations, for example message or content filtering, DLP, or disclaimer stamping. Attachment names are available only for messages that pass through body scanning while the attachment is still attached. Some examples when an attachment name will not appear include (but are not limited to):

  - if the system only uses content filters, and a message is dropped or its attachment is stripped by anti-spam or anti-virus filters
  - if message splintering policies strip the attachment from some messages before body scanning occurs.

- **File SHA256**: Find messages with the SHA-256 value of the message file
- **Cisco Host**: Narrow search criteria to particular Email Security appliances, or search across all managed appliances.
- **Message ID Header and Cisco MID**: Find messages by identifying the SMTP “Message-ID:” header or the Cisco message ID (MID).
- **Sender IP Address/ Domain/ Network Owner**: Search for messages from a particular IP address, domain name or network owner.
- **Message Event**: Find messages that match specified events, such as messages flagged as virus positive, spam positive, or suspected spam, and messages that were delivered, hard bounced, soft bounced, or sent to the Virus Outbreak Quarantine
- **Rejected Connections**: Search for messages from a particular IP address, domain name or network owner of the rejected connections in the search results

Setting Up Centralized Message Tracking

To set up centralized message tracking, complete the following procedures in order:

Enabling Centralized Email Tracking on a Security Management Appliance

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Email> Centralized Message Tracking**.

Step 3 In the Message Tracking Service section, click **Enable**.

Step 4 If you are enabling centralized email tracking for the first time after running the System Setup Wizard, review the end user license agreement, and click **Accept**.

Step 5 **Submit** and commit your changes.

Configuring Centralized Message Tracking on Email Security Appliances

Step 1 Verify that Message Tracking is configured and working properly on the Email Security appliance.

Step 2 Go to **Security Services > Message Tracking**.

Step 3 Click **Edit Settings**.

Step 4 Select **Centralized Tracking**.

Step 5 Click **Submit**.

Step 6 If you want to be able to search for and log the names of email attachments:
Make sure you have at least one incoming content filter or other body scanning feature configured and enabled on the Email Security appliance. For information about content filters and body scanning, see the documentation or online help for your Email Security appliance.

Step 7 Commit your changes.

Step 8 Repeat for each Email Security appliance to manage.

Adding the Centralized Message Tracking Service to Each Managed Email Security Appliance

The steps you follow depend on whether or not you have already added the appliance while configuring another centralized management feature.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Security Appliances**.

Step 3 If you have already added the Email Security appliance to the list on this page:

a) Click the name of an Email Security appliance.
b) Select the **Centralized Message Tracking** service.

Step 4 If you have not yet added the Email Security appliance:

a) Click Add Email Appliance.
b) In the Appliance Name and IP Address text fields, type the appliance name and the IP address for the Management interface of the Email Security appliance.
Note If you enter a DNS name in the IP Address text field, it will be immediately resolved to an IP address when you click **Submit**.
c) The Centralized Message Tracking service is pre-selected.
d) Click **Establish Connection.**
e) Enter the user name and password for an administrator account on the appliance to be managed, then click **Establish Connection**.
Note You enter the login credentials to pass a public SSH key for file transfers from the Security Management appliance to the remote appliance. The login credentials are not stored on the Security Management appliance.
f) Wait for the Success message to appear above the table on the page.
g) Click **Test Connection**.
h) Read test results above the table.

Step 5 Click **Submit**.

Step 6 Repeat this procedure for each Email Security appliance for which you want to enable Centralized Message Tracking.

Step 7 Commit your changes.

Managing Access to Sensitive Information

If you will distribute administrative tasks to other people and you want to restrict their access to sensitive information that may appear in email messages that violate Data Loss Prevention (DLP) policies, see Controlling Access to Sensitive Information in Message Tracking , on page 276.

Checking Message Tracking Data Availability

You can determine the date range that your message tracking data includes, as well as identify any missing intervals in that data.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Message Tracking > Message Tracking Data Availability**.

Searching for Email Messages

Searching for Email Messages on the Cloud Email Security Management Console

The tracking service of the appliance lets you search for a particular email message or group of messages that match specified criteria, such as the message subject line, date and time range, envelope sender or recipient, or processing event (for example, whether the message was virus positive, spam positive, hard bounced, delivered, and so forth). Message tracking gives you a detailed view of message flow. You can also drill down on particular email messages to see message details, such as the processing events, attachment names, or the envelope and header information.

Note Although the tracking component provides detailed information about individual email messages, you cannot use it to read the content of messages.

Step 1 On the Cloud Email Security Management Console, choose **Tracking** > **Search**.

Step 2 Select **Messages** tab or **Rejected Connections** tab to narrow your search results.

Note You can search for rejected connections based on the sender IP address, domain or network owner.

Step 3 (Optional) Click the **Advanced Search** to display additional search options.

Step 4 Enter the following search criteria:

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

- [For Messages and Rejected Connections] **Message Received**: Specify a date and time range for the query using “Last Day,” “Last 7 Days,” or “Custom Range.” Use the “Last Day” option to search for messages within the past 24 hours, and use the “Last 7 Days” option to search for messages within the past full seven days, plus the time that has passed on the current day.

If you do not specify a date, the query returns data for all dates. If you specify a time range only, the query returns data for that time range across all available dates. If you specify the current date and 23:59 as the end date and time, the query returns all data for the current date.

Dates and times are converted to GMT format when they are stored in the database. When you view dates and times on an appliance, they are displayed in the local time of the appliance.

Messages appear in the results only after they have been logged on the Email Security appliance and retrieved by the Security Management appliance. Depending on the size of logs and the frequency of polling, there could be a small gap between the time when an email message was sent and when it actually appears in tracking and reporting results.

- **Envelope Sender**: Select Begins With, Is, or Contains, and enter a text string to search for in the envelope sender. You can enter email addresses, user names, or domains. Use the following formats:
  - For email domains: *example.com, [203.0.113.15], [ipv6:2001:db8:80:1::5]*
  - For full email addresses: *user@example.com, user@[203.0.113.15] or user@[ipv6:2001:db8:80:1::5]*.
  - You can enter any character(s). No validation of your entry is performed.
- **Subject**: Select Begins With, Is, Contains, or Is Empty, and enter a text string to search for in the message subject line.
- **Envelope Recipient**: Select Begins With, Is, or Contains, and enter text to search for in the envelope recipient. You can enter email addresses, user names, or domains.

  If you use the alias table for alias expansion on your Email Security appliances, the search finds the expanded recipient addresses rather than the original envelope addresses. In all other cases, message tracking queries find the original envelope recipient addresses.

  Otherwise, valid search criteria for Envelope Recipient are the same as those for Envelope Sender.

  You can enter any character(s). No validation of your entry is performed.
- **Attachment Name**: Select Begins With, Is, or Contains, and enter an ASCII or Unicode text string for one Attachment Name to find. Leading and trailing spaces are not stripped from the text you enter.
- **File SHA256**: Enter a File SHA-256 value of the message.

  For more information about identifying files based on SHA-256 hash, see Identifying Files by SHA-256 Hash , on page 67.
- **Cisco Host**: Select All Host to search across all email security appliances or select the required email security appliance from the drop-down menu.
- [For Messages and Rejected Connections] **Show Additional Details (Message Last State, SBRS, Sender IP, and Group)**: Select this check box to show all the additional details such as verdict chart, last state of the message, SBRS score, sender IP address and sender groups.

  Note Depending on the size of your database, it may take a few minutes to display the search results.
- **Message ID Header and Cisco MID**: Enter a text string for the message ID header, the Cisco IronPort message ID (MID), or both.
- [For Messages and Rejected Connections] **Sender IP Address/ Domain/ Network Owner**: Enter a sender IP address, domain or nework owner details.

- An IPv4 address must be 4 numbers separated by a period. Each number must be a value from 0 to 255. (Example: 203.0.113.15).
- An IPv6 address consists of 8 sets of 16-bit hexadecimal values separated by colons.

  You can use zero compression in one location, such as 2001:db8:80:1::5.
- **Message Event**: Select the events to track. Options are Virus Positive, Spam Positive, Suspect Spam, contained malicious URLs, contained URL in specified category, DLP Violations (you can enter the name of a DLP policy and select violation severities or action taken), DMARC violations, Delivered, Advanced Malware Protection Positive (for malware found in an attachment), Hard Bounced, Soft Bounced, currently in a policy, virus, or outbreak quarantine, caught by message filters or content filters, and Quarantined as Spam. Unlike most conditions that you add to a tracking query, events are added with an "OR" operator. Selecting multiple events expands the search.

You do not need to complete every field. Except for the Message Event options, the query is an "AND" search. The query returns messages that match the "AND" conditions specified in the search fields. For example, if you specify text strings for the envelope recipient and the subject line parameters, the query returns only messages that match both the specified envelope recipient and the subject line.

Step 5 Click **Search**.
Each row corresponds to an email message. Scroll down to load more messages in the view.

If necessary, you can refine your search by entering new search criteria, and run the query again. Alternatively, you can refine the search by narrowing the result set, as described in the following section.

Searching for Email Messages on the Legacy Web Interface

The Security Management appliance's tracking service lets you search for a particular email message or group of messages that match specified criteria, such as the message subject line, date and time range, envelope sender or recipient, or processing event (for example, whether the message was virus positive, spam positive, hard bounced, delivered, and so forth). Message tracking gives you a detailed view of message flow. You can also drill down on particular email messages to see message details, such as the processing events, attachment names, or the envelope and header information.

Note Although the tracking component provides detailed information about individual email messages, you cannot use it to read the content of messages.

Step 1 Choose **Email > Message Tracking > Message Tracking**.

Step 2 (Optional) Click the Advanced link to display more search options.

Step 3 Enter search criteria:

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

- Envelope Sender: Select Begins With, Is, or Contains, and enter a text string to search for in the envelope sender. You can enter email addresses, user names, or domains. Use the following formats:

- For email domains: example.com, [203.0.113.15], [ipv6:2001:db8:80:1::5]
- For full email addresses: user@example.com, user@[203.0.113.15] or user@[ipv6:2001:db8:80:1::5].
- You can enter any character(s). No validation of your entry is performed.

- Envelope Recipient: Select Begins With, Is, or Contains, and enter text to search for in the envelope recipient. You can enter email addresses, user names, or domains.

If you use the alias table for alias expansion on your Email Security appliances, the search finds the expanded recipient addresses rather than the original envelope addresses. In all other cases, message tracking queries find the original envelope recipient addresses.

Otherwise, valid search criteria for Envelope Recipient are the same as those for Envelope Sender.

You can enter any character(s). No validation of your entry is performed.

- Subject: Select Begins With, Is, Contains, or Is Empty, and enter a text string to search for in the message subject line.
- Message Received: Specify a date and time range for the query using “Last Day,” “Last 7 Days,” or “Custom Range.” Use the “Last Day” option to search for messages within the past 24 hours, and use the “Last 7 Days” option to search for messages within the past full seven days, plus the time that has passed on the current day.

If you do not specify a date, the query returns data for all dates. If you specify a time range only, the query returns data for that time range across all available dates. If you specify the current date and 23:59 as the end date and time, the query returns all data for the current date.

Dates and times are converted to GMT format when they are stored in the database. When you view dates and times on an appliance, they are displayed in the local time of the appliance.

Messages appear in the results only after they have been logged on the Email Security appliance and retrieved by the Security Management appliance. Depending on the size of logs and the frequency of polling, there could be a small gap between the time when an email message was sent and when it actually appears in tracking and reporting results.

- Sender IP Address: Enter a sender IP address and select whether to search messages or to search rejected connections only.
  - An IPv4 address must be 4 numbers separated by a period. Each number must be a value from 0 to 255. (Example: 203.0.113.15).
  - An IPv6 address consists of 8 sets of 16-bit hexadecimal values separated by colons. You can use zero compression in one location, such as 2001:db8:80:1::5.
- Message ID Header and Cisco IronPort MID: Enter a text string for the message ID header, the Cisco IronPort message ID (MID), or both.
- Query Settings: From the drop-down menu, select how long you want the query to run before it times out. Options are “1 minute,” “2 minutes,” “5 minutes,” “10 minutes,” and “No time limit.” Also, select the maximum number of results you want the query to return (up to 1000).
- Attachment name: Select Begins With, Is, or Contains, and enter an ASCII or Unicode text string for one Attachment Name to find. Leading and trailing spaces are not stripped from the text you enter.

You do not need to complete every field. Except for the Message Event options, the query is an “AND” search. The query returns messages that match the “AND” conditions specified in the search fields. For example, if you specify text strings

for the envelope recipient and the subject line parameters, the query returns only messages that match both the specified envelope recipient and the subject line.

Step 4 Click **Search**.
The query results appear at the bottom of the page. Each row corresponds to an email message.

Your search criteria are highlighted in each row.

If the number of returned rows is greater than the value specified in the "Items per page" field, the results appear on multiple pages. To navigate through the pages, click the page numbers at the top or bottom of the list.

If necessary, refine the search by entering new search criteria, and run the query again. Alternatively, you can refine the search by narrowing the result set, as described in the following section.

Narrowing the Result Set

After you run a query, you might find that the result set includes more information than you need. Instead of creating a new query, narrow the result set by clicking a value within a row in the list of results. Clicking a value adds the parameter value as a condition in the search. For example, if the query results include messages from multiple dates, click a particular date within a row to show only messages that were received on that date.

Step 1 Float the cursor over the value that you want to add as a condition. The value is highlighted in yellow.
Use the following parameter values to refine the search:

- Date and time
- Message ID (MID)
- Host (the Email Security appliance)
- Sender
- Recipient
- The subject line of the message, or starting words of the subject

Step 2 [New Web Interface Only] In the Message Tracking search criteria, click **Modify**.
Use the following parameter values to refine the search:

- Date and time
- Message ID (MID)
- Cisco Host (the Email Security appliance)
- Sender
- Recipient
- The subject line of the message, or starting words of the subject
- Message Event

- Additional Details (Message Last State, SBRS, Sender IP, and Group)

Step 3 Click the value to refine the search.
The Results section displays the messages that match the original query parameters *and* the new condition that you added.

Step 4 If necessary, click additional values in the results to further refine the search.
Note To remove query conditions, click **Clear** and run a new tracking query.

About Message Tracking and Advanced Malware Protection Features

When searching for file threat information in Message Tracking, keep the following points in mind:

- To search for malicious files found by the file reputation service, select **Advanced Malware Protection Positive** for the Message Event option in the Advanced section in Message Tracking.
- Message Tracking includes only information about file reputation processing and the original file reputation verdicts returned at the time a message was processed. For example, if a file was initially found to be clean, then a verdict update found the file to be malicious, only the clean verdict appears in Tracking results.

  In Message Tracking details, the Processing Details section shows:

  - The SHA-256 of each attachment in the message, and
  - The final Advanced Malware Protection verdict for the message as a whole, and
  - Any attachments which were found to contain malware.

    No information is provided for clean or unscannable attachments.

- Verdict updates are available only in the AMP Verdict Updates report. The original message details in Message Tracking are not updated with verdict changes. To see messages that have a particular attachment, click a SHA-256 in the verdict updates report.
- Information about File Analysis, including analysis results and whether or not a file was sent for analysis, are available only in the File Analysis report.

  Additional information about an analyzed file may be available from the cloud. To view any available File Analysis information for a file, select **Monitor > File Analysis** and enter the SHA-256 to search for the file. If the File Analysis service has analyzed the file from any source, you can see the details. Results are displayed only for files that have been analyzed.

  If the appliance processed a subsequent instance of a file that was sent for analysis, those instances will appear in Message Tracking search results.

Understanding Tracking Query Results

If results are not what you expected, see Troubleshooting Message Tracking, on page 172.

Tracking query results list all of the messages that match the criteria specified in the tracking query. Except for the Message Event options, the query conditions are added with an “AND” operator. The messages in the

result set must satisfy all of the “AND” conditions. For example, if you specify that the envelope sender begins with J and you specify that the subject begins with T , the query returns a message only if both conditions are true for that message.

To view detailed information about a message, click the **More Details** link in the new web interface or **Show Details** linkin the legacy web interface, for that message. For more information, see the Message Details, on page 169.

Note

- Messages with 50 or more recipients will not appear in tracking query results. This issue will be resolved in a future release.
- [New Web Interface Only] You can scroll down to display the search results when you specify your query. More results are displayed in the view as you scroll down.
- You can export the search results to a .csv file using the **Export** link above the search results section.

  You can choose to display up to 1000 search results when you specify your query. To view up to 50,000 messages that match your criteria, click the **Export All** link above the search results section and open the resulting .csv file in another application.
- If you clicked a link in a report page to view message details in Message Tracking, and the results are unexpected, this can occur if reporting and tracking were not both simultaneously and continuously enabled during the time period you are reviewing.
- For information about printing or exporting message tracking search results, see Printing and Exporting Reporting Data , on page 31.

Message Details

To view detailed information about a particular email message, including the message header information and processing details, click **More Details** linkfor any item in the search results list. A new window opens with the message details.

The message details include the following sections:

Verdict Chart and Last State Verdicts

A Verdict Chart displays information of the various possible verdicts triggered by each engine of the email security appliance.

The following figure shows the various verdicts of each engine:

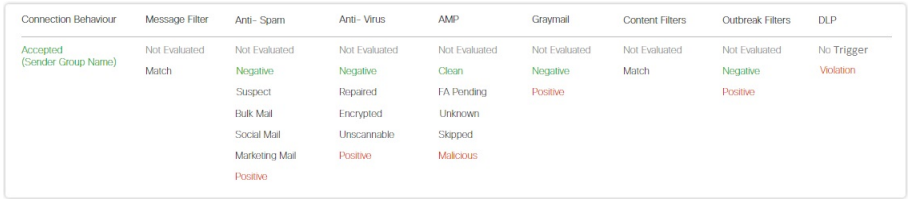

Note The service engines that are disabled or license expired are displayed with a verdict “Not Applicable”.

The Last State verdict of the message determines the final verdict that is triggered after all the possible verdicts of each of the engine in your appliance.

Following are some of the last state verdicts:

- **Delivered**: When a message is delivered.
- **Dropped**: When a message is dropped.
- **Bounced**: When a message is bounced back.
- **Splintered**: When the MID of a message is split into multiple MIDs having multiple final states.
- **Quarntined**: When a message is quarantined by engines.
- **Queued**: When a message is still being processed by an engine and gives out an interim verdict; or when a message is waiting in a queue of a particular engine.
- **Processing**: When a message is not completely processed by all the engines; or when a message is waiting in a queue of a particular engine.
- **Last State Not Available**: When both last state and verdict charts cannot be retrived, you can view the Last State with no drop-down.

Envelope and Header Summary

This section displays information from the message envelope and header, such as the envelope sender and recipients. It includes the following information:

Received Time: Time that the Email Security appliance received the message.

MID: Message ID.

Subject: Subject line of the message.

The subject line in the tracking results may have the value “(No Subject)” if the message does not have a subject or if the Email Security appliances are not configured to record the subject lines in log files.

Envelope Sender: Address of the sender in the SMTP envelope.

Envelope Recipients: Addresses of the recipients in the SMTP envelope.

Message ID Header: “Message-ID:” header that uniquely identifies each email message. It is inserted in the message when the message is first created. The “Message-ID:” header can be useful when you are searching for a particular message.

Cisco Host: Email Security appliance that processed the message.

SMTP Auth User ID: SMTP authenticated user name of the sender, if the sender used SMTP authentication to send the email. Otherwise, the value is “N/A.”

Attachments: The names of files attached to the message.

Sender Group: The sender group that received the message.

Message Size: The size of the message.

Policy Match (Incoming or Outgoing): The policy that received the message.

Note If the engine is not able to fetch the details, the value is displayed as “N/A”.

Sending Host Summary

Reverse DNS Hostname: Hostname of the sending host, as verified by reverse DNS (PTR) lookup.

IP Address: IP address of the sending host.

SBRS Score: (SenderBase Reputation Score). The range is from 10 (likely a trustworthy sender) to -10 (apparent spammer). A score of "None" indicates that there was no information about this host at the time the message was processed.

Processing Details

This section displays various logged status events during the processing of the message.

Entries include information about mail policy processing, such as anti-spam and anti-virus scanning, and other events such as message splitting.

If the message was delivered, the details of the delivery appear here. For example, a message may have been delivered and a copy kept in quarantine.

The last recorded event is highlighted in the processing details.

Summary Tab

This tab displays the summary logs of all the events during the processing of message.

DLP Matched Content Tab

This tab displays content that violates Data Loss Prevention (DLP) policies.

Because this content typically includes sensitive information, such as corporate confidential information or personal information including credit card numbers and health records, you may want to disable access to this content for users who have access, but not Administrator-level access, to the Security Management appliance. See Controlling Access to Sensitive Information in Message Tracking , on page 276.

URL Details Tab

This tab displays only for messages caught by URL Reputation and URL Category content filters and by outbreak filters not message filters.

This tab displays the following information:

- The reputation score or category associated with the URL
- The action performed on the URL (rewrite, defang, or redirect)
- If a message contains multiple URLs, which URL has triggered the filter action.

You can see this tab only if you have configured your Email Security appliance to display this information. See *User Guide for AsyncOS for Cisco Email Security Appliances* .

To control access to this tab, see Controlling Access to Sensitive Information in Message Tracking , on page 276

SMTP Log Tab

This section displays a log of messages when the sender of the email fails SMTP authentication.

AMP Log Tab

This section displays a log of messages caught by the Advanced Malware Protection file reputation and file analysis service.

Troubleshooting Message Tracking

Expected Messages Are Missing from Search Results

Problem

Search results did not include messages that should have met the criteria.

Solution

- Results for many searches, especially Message Event searches, depend on your appliance configuration. For example, if you search for a URL Category for which you have not filtered, no results will be found, even if a message contains a URL in that category. Verify that you have configured the Email Security appliance properly to achieve the behavior that you expected. For example, check your mail policies, content and message filters, and quarantine settings.
- See Checking Message Tracking Data Availability , on page 163.

Attachments Do Not Appear in Search Results

Problem

Attachment names are not found and displayed in search results.

Solution

At least one incoming content filter or other body scanning feature configured and enabled on the ESA. See configuration requirements at Enabling Centralized Email Tracking on a Security Management Appliance, on page 161 and limitations for attachment name searches in Tracking Service Overview, on page 159.

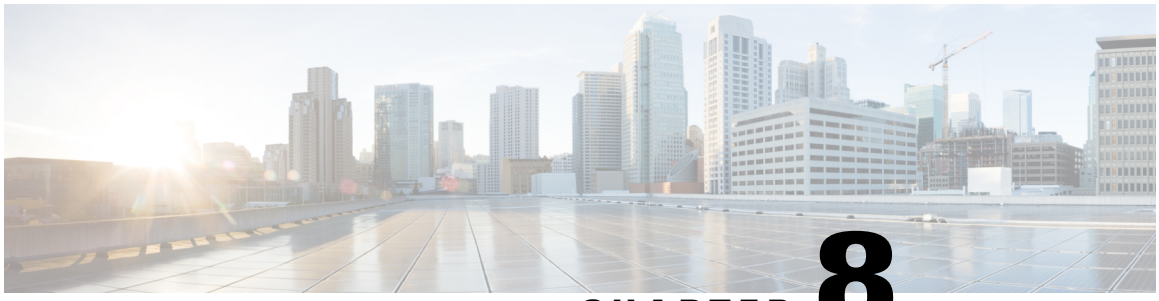

CHAPTER 8

Spam Quarantine

This chapter contains the following sections:

Overview of the Spam Quarantine

The spam quarantine (also known as ISQ, End-User Quarantine, and EUQ) provides a safeguard mechanism for organizations that are concerned about "false positives" — that is, legitimate email messages that the appliance has deemed to be spam. When the appliance determines that a message is spam or suspected spam, you may want to let the recipient or an administrator review the message before delivering or deleting it. The spam quarantine stores messages for this purpose.

Administrative users of the Email Security appliance can view all messages in a spam quarantine. End users, usually the message recipients, can view their own quarantined messages in a slightly different web interface.

The spam quarantine is separate from policy, virus, and outbreak quarantines.

Local Versus External Spam Quarantine

A local spam quarantine stores spam and suspect spam on the Email Security appliance. An external spam quarantine can store these messages on a separate Cisco Content Security Management appliance.

Consider using an external spam quarantine if:

- You want a centralized location to store and manage spam from multiple Email Security appliances.
- You want to store more spam than the Email Security appliance can hold.
- You want to regularly back up the spam quarantine and its messages.

Setting Up the Centralized Spam Quarantine

DETAILED STEPS

	Command or Action	Purpose
Step 1	On the Security Management appliance, enable the centralized spam quarantine service.	Enabling and Configuring the Spam Quarantine , on page 175
Step 2	On the Security Management appliance, specify the Email Security appliances to include in the centralized spam quarantine.	Adding the Centralized Spam Quarantine Service to Each Managed Email Security Appliance , on page 177
Step 3	Set up the Security Management appliance for sending notifications and released spam.	Configuring an Outbound IP Interface on the Security Management Appliance , on page 177
Step 4	On the Security Management appliance, configure the spam quarantine browser interface.	Configuring the IP Interface for Browser Access to the Spam Quarantine , on page 178
Step 5	Ensure that the Email Security appliance is configured to send mail to the spam quarantine.	In the documentation for your Email Security appliance, see information about configuring anti-spam and mail policies. Links to relevant sections are in the table in the section on setting up a local spam quarantine.
Step 6	On the Email Security appliances, enable and configure the external spam quarantine.	See the documentation for your Email Security appliance.
Step 7	On the Email Security appliances, disable the local quarantine.	See information about disabling the local spam quarantine to activate the external spam quarantine in the documentation for your Email Security appliance.

Enabling and Configuring the Spam Quarantine

Note If you use an external spam quarantine, you will configure the settings described in this section on the Security Management appliance.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Spam Quarantine**.

Step 3 If you are enabling the spam quarantine for the first time after running the System Setup Wizard:

a) Click **Enable**.

b) Review the end user license agreement, then click **Accept**.

Step 4 If you are editing spam quarantine settings, click **Edit Settings**.

Step 5 Specify options:

Option	Description
Quarantine IP Interface Quarantine Port	By default, the spam quarantine uses the Management interface and port 6025. The IP interface is the interface on the Security Management appliance that is configured to listen for incoming mail. The quarantine port is the port number that the sending appliances use in their external quarantine settings. If your Email Security appliances are not on the same network as your Security Management appliance, then you must use the Management interface.
Deliver Messages Via	All outgoing quarantine-related email (such as spam notifications and messages released from the spam quarantine) must be delivered via another appliance or server that is configured to send messages. You can route these messages through an SMTP or groupware server, or you can specify the outbound listener interface of an Email Security appliance (typically the Data 2 interface). The alternate address is used for load balancing and failover. If you have multiple Email Security appliances, you can use the outbound listener interface of any managed Email Security appliances for the primary and alternate addresses. Both must use the same interface (either Data 1 or Data 2) as the outbound listener. Read instructions on the screen for additional caveats about these addresses.

Option	Description
Quarantine Size	If you deselect **When storage space is full, automatically delete oldest messages first**, newer messages will not be added to a full quarantine. Cisco recommends that you enable this option so that a full quarantine will not cause messages to queue (back up) on your appliance. To manage disk space for your quarantine, see Managing Disk Space , on page 333 .
Schedule Delete After	Specify the number of days to hold messages before deleting them. Cisco recommends that you configure the quarantine to delete older messages to prevent the quarantine from filling to capacity, but you can elect not to schedule automatic deletion.
Notify Cisco Upon Message Release	—
Spam Quarantine Appearance	**Logo** By default, the Cisco logo is displayed at the top of the spam quarantine page when the user logs in to view quarantined messages. To use a custom logo instead, upload the logo. The logo should be a .jpg, .gif, or .png file that is at most 50 pixels high by 500 pixels wide.
	Login page message (Optional) Specify a login page message. This message is shown to end users and administrators when they log in to view the quarantine. If you do not specify a message, the following message appears: Enter your login information below. If you are unsure what to enter, please contact your administrator.
Administrative Users	See Configuring Administrative User Access to the Spam Quarantine , on page 179.

Step 6 Submit and commit your changes.

What to Do Next

- Return to Setting Up the Centralized Spam Quarantine , on page 174

Adding the Centralized Spam Quarantine Service to Each Managed Email Security Appliance

The steps that you follow depend on whether or not you have already added the appliance while configuring another centralized management feature.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Security Appliances**.

Step 3 If you have already added the Email Security appliance to the list on this page:

a) Click the name of the Email Security appliance.
b) Select the **Spam Quarantine** service.

Step 4 If you have not yet added Email Security appliances:

a) Click Add Email Appliance.
b) In the Appliance Name and IP Address text fields, type the appliance name and the IP address for the Management interface of the appliance.
Note A DNS name may be entered in the IP Address text field; however, it will be immediately resolved to an IP address when you click **Submit**.
c) The Spam Quarantine service is pre-selected.
d) Click **Establish Connection.**
e) Enter the username and password for an administrator account on the appliance to be managed, then click **Establish Connection**.
Note You enter the login credentials to pass a public SSH key for file transfers from the Security Management appliance to the remote appliance. The login credentials are not stored on the Security Management appliance.
f) Wait for the success message to appear above the table on the page.
g) Click **Test Connection**.
h) Read the test results above the table.

Step 5 Click **Submit**.

Step 6 Repeat this procedure for each Email Security appliance for which you want to enable the spam quarantine.

Step 7 Commit your changes.

Configuring an Outbound IP Interface on the Security Management Appliance

Configure an interface on the Security Management appliance to send quarantine-related messages (including notifications and released email) to the Email Security appliance for delivery.

Before You Begin

Obtain or identify an IP address to use for the outbound interface. This will typically be the Data 2 interface on the Security Management appliance. For more information about network requirements, see Assigning Network and IP Addresses, on page 397

Step 1 Use this procedure in conjunction with the information in Configuring IP Interfaces, on page 390

Step 2 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 3 Choose **Management Appliance > Network IP Interfaces**.

Step 4 Click **Add IP Interface**.

Step 5 Enter the following settings:

- Name
- Ethernet Port

  Typically, this will be Data 2. Specifically, this must match the data interface on the Email Security appliance that you specified for the **Primary Server** in the **Deliver Messages Via** section of the Spam Quarantine Settings page under **Management Appliance > Centralized Services > Spam Quarantine**.
- IP Address

  IP address of the interface that you just specified.
- Netmask
- Hostname

  For example, if this is the Data 2 interface, use data2.sma.example.com .

Do not enter information in the Spam Quarantine section for this interface.

Step 6 Submit and commit your changes.

Configuring the IP Interface for Browser Access to the Spam Quarantine

When administrators and end users access the spam quarantine, a separate browser window opens.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Network > IP Interfaces**.

Step 3 Click the name of the **Management** interface.

Step 4

Step 5 In the Spam Quarantine section, configure settings for access to the spam quarantine:

- By default, HTTP uses port 82 and HTTPS uses port 83.

- Specify the URL that appears in notifications and in the spam quarantine browser window.

  If you do not want to expose the hostname of your Security Management appliance to end users, you can specify an alternate hostname.

Step 6 Submit and commit your changes.

What to Do Next

Ensure that your DNS server can resolve the hostname that you specified for spam quarantine access.

Configuring Administrative User Access to the Spam Quarantine

All users with administrator privileges can change spam quarantine settings and view and manage messages in the spam quarantine. You do not need to configure spam quarantine access for administrator users.

If you configure access to the spam quarantine for users with the following roles, they can view, release, and delete messages in the spam quarantine:

- Email administrator
- Operator
- Read-only operator
- Help desk user
- Guest
- Custom user roles that have spam quarantine privileges

These users cannot access spam quarantine settings.

Before You Begin

Create users or custom user roles that have access to the spam quarantine. For more information, see information about Access to Quarantines for Custom User Role , on page 258 in Distributing Administrative Tasks, on page 253

Step 1 If you are not already editing the spam quarantine settings page:

a) [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

b) Choose **Management Appliance > Centralized Services** > **Spam Quarantine**.

c) Click **Edit Settings** the **Spam Quarantine** link in the Quarantine Name column of the Spam Quarantine section.

Step 2 Click the link for the type of user to add: local, externally authenticated, or custom role.
If you have already added users or roles, click a username or role to view all eligible users or roles.

Step 3 Select the users or roles that you want to add.
Users with Administrator privileges , including Email Administrators, are not listed because they automatically have full access to the spam quarantine.

Step 4 Click **OK**.

Step 5 Submit and commit your changes.

Limiting Which Recipients Have Mail Quarantined

You can use multiple mail policies(Mail Policies > Incoming Mail Policy) to specify a list of recipient addresses for which mail will not be quarantined. Select 'Deliver' or 'Drop' instead of quarantine when configuring the anti-spam settings for the mail policy.

Spam Quarantine Language

Each user selects a language in the spam quarantine from the Options menu at the top right of the window.

Edit Spam Quarantine Page

Using Safelists and Blocklists to Control Email Delivery Based on Sender

Administrators and end users can use safelists and blocklists to help determine which messages are spam. Safelists specify senders and domains that are never treated as spam. Blocklists specify senders and domains that are always treated as spam.

You can allow end users (email users) to manage the safelist and blocklist for their own email accounts. For example, an end user may receive email from a mailing list that no longer interests him. He may decide to add this sender to his blocklist to prevent emails from the mailing list from being sent to his inbox. On the other hand, end users may find that emails from specific senders are sent to their spam quarantine when they do not want them to be treated as spam. To ensure that messages from these senders are not quarantined, they may want to add the senders to their safelists.

Changes that end users and administrators make are visible to and can be changed by either.

Message Processing of Safelists and Blocklists

A sender's being on a safelist or blocklist does not prevent the appliance from scanning a message for viruses or determining if the message meets the criteria for a content-related mail policy. Even if the sender of a message is on the recipient's safelist, the message may not be delivered to the end user depending on other scanning settings and results.

When you enable safelists and blocklists, the appliance scans the messages against the safelist/blocklist database immediately before anti-spam scanning. If the appliance detects a sender or domain that matches a safelist or blocklist entry, the message will be splintered if there are multiple recipients (and the recipients have different safelist/blocklist settings). For example, a message is sent to both recipient A and recipient B. Recipient A has safelisted the sender, whereas recipient B does not have an entry for the sender in the safelist or the blocklist. In this case, the message may be split into two messages with two message IDs. The message

sent to recipient A is marked as safelisted with an *X-SLBL-Result-Safelist* header and skips anti-spam scanning, whereas the message bound for recipient B is scanned by the anti-spam scanning engine. Both messages then continue along the pipeline (through anti-virus scanning, content policies, and so on) and are subject to any configured settings.

If a message sender or domain is blocklisted, the delivery behavior depends on the blocklist action that you specify when you enable the safelist/blocklist feature. Similar to safelist delivery, the message is splintered if there are different recipients with different safelist/blocklist settings. The blocklisted message splinter is then quarantined or dropped, depending on the blocklist action settings. If the blocklist action is configured to quarantine, the message is scanned and eventually quarantined. If the blocklist action is configured to delete, the message is dropped immediately after safelist/blocklist scanning.

Because safelists and blocklists are maintained in the spam quarantine, delivery behavior is also contingent on other anti-spam settings. For example, if you configure the "Accept" mail flow policy in the Host Access Table (HAT) to skip anti-spam scanning, then users who receive mail on that listener will not have their safelist and blocklist settings applied to mail received on that listener. Similarly, if you create a mail flow policy that skips anti-spam scanning for certain message recipients, these recipients will not have their safelist and blocklist settings applied.

Enabling Safelists and Blocklists

Before You Begin

- The spam quarantine must be enabled. See Setting Up the Centralized Spam Quarantine , on page 174.
- Configure the Email Security appliance to use an external safelist/blocklist. See instructions for setting up an external spam quarantine in the documentation for your Email Security appliance.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Spam Quarantine**.

Step 3 In the **End-User Safelist/Blocklist (Spam Quarantine)** section, select **Enable**.

Step 4 Select **Enable End User Safelist/Blocklist Feature**.

Step 5 Specify the **Maximum List Items Per User**.
This is the maximum number of addresses or domains for each list, for each recipient. If you allow a large number of list entries per user, system performance might be adversely affected.

Step 6 Submit and commit your changes.

External Spam Quarantine and Safelist/Blocklists

Because the Email Security appliance evaluates senders in safelists and blocklists when processing incoming mail, safelists and blocklists that are stored on a Security Management appliance must be sent to the Email Security appliance in order to be applied to incoming mail. When you configure the safelist/blocklist feature on a Security Management appliance, you configure the frequency of these updates.

Adding Senders and Domains to Safelists and Blocklists (Administrators)

Manage safelists and blocklists via the spam quarantine interface.

You can also see whether many recipients (end users in your organization) have whitelisted or blacklisted a particular sender or domain.

Administrators see and work with the superset of the same entries that each end user sees and works with.

Before You Begin

- Ensure that you can access the spam quarantine. See Accessing the Spam Quarantine (Administrative Users), on page 194.
- Enable access to the safelist/blocklist. See Enabling Safelists and Blocklists , on page 181.
- (Optional) To import safelist/blocklists instead of building these lists using the procedure in this section, use the process described in Backing Up and Restoring the Safelist/Blocklist , on page 185.
- Understand the required format of safelist and blocklist entries. See Syntax for Safelists and Blocklist Entries , on page 183.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click **Quarantine** > **Spam Quarantine**. You must login to the legacy web interface again.

Step 2 Click on the spam quarantine link.

Step 3 Select the **Options** drop-down menu in the upper right corner of the page.

Step 4 Choose **Safelist** or **Blocklist**.

Step 5 (Optional) Search for a sender or recipient.

Step 6 Do one or more of the following:

To	Do This
Add multiple senders for a recipient	1 Select **View by: Recipient** 2 Click **Add**, or click **Edit** for a recipient. 3 Enter or edit the recipient email address. 4 Enter sender email addresses and domains. Put each entry on a separate line, or separate each entry with a comma. 5 Click **Submit**.

To	Do This
Add multiple recipients for a sender	1 Select **View by: Sender** 2 Click **Add**, or click **Edit** for a sender. 3 Enter or edit the sender address or domain. 4 Enter recipient email addresses. Put each entry on a separate line, or separate each entry with a comma. 5 Click **Submit**.
Delete all senders associated with a recipient Delete all recipients associated with a sender	1 Select a **View by** option. 2 Click a trash can icon to delete an entire table row.
Delete individual senders for a recipient Delete individual recipients for a sender	1 Select a View by option. 2 Click **Edit** for an individual recipient or sender. 3 Add or remove entries from the text box. You must leave at least one entry. 4 Click **Submit**.

Syntax for Safelists and Blocklist Entries

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

- user@domain.com
- server.domain.com
- domain.com
- [10.1.1.0]
- [ipv6:2001:DB8:1::1]
- user@[1.2.3.4]
- user@[ipv6:2001:db8::1]

An identical entry, such as a sender address or a domain, cannot be included on both the safelist and the blocklist at the same time. However, a domain can be on a safelist while an email address for a sender belonging to that domain is on the blocklist (or vice versa), and both rules apply. For example, if *example.com* is on the safelist, *george@example.com* can be on the blocklist. In this case, the appliance delivers all mail from *example.com* without scanning for spam, except mail from *george@example.com,* which is treated as spam.

It is not possible allow or block a range of subdomains using the following syntax: *.domain.com* . However, it is possible to block a specific domain using the following syntax: *server.domain.com* .

Clearing All Safelists and Blocklists

If you need to delete all safelist and blocklist entries, including all senders and all recipients, import a file with no entries using the procedure in Backing Up and Restoring the Safelist/Blocklist , on page 185.

About End-User Access to Safelists and Blocklists

End users access their safelist and blocklist via the spam quarantine. To configure end-user access to the spam quarantine, see Setting Up End-User Access to the Spam Quarantine via Web Browser , on page 189.

You may want to give your end users the URL of the spam quarantine and the instructions below, as applicable.

Adding Entries to Safelists (End Users)

Note Delivery of messages from safelisted senders depends on other settings that are configured in the system. See Message Processing of Safelists and Blocklists , on page 180.

End users can add senders to safelists in two ways:

Adding the Sender of a Quarantined Message to the Safelist

End users can add senders to the safelist if the message has been sent to the spam quarantine.

Step 1 From the spam quarantine, select the checkbox next to the message.

Step 2 Choose **Release and Add to Safelist** from the drop-down menu.
The envelope sender and the from header for the specified mail are both added to the safelist, and the released messages proceed directly to the destination queue, skipping any further work queue processing in the email pipeline.

Adding Senders to the Safelist Without a Quarantined Message

Step 1 Access the spam quarantine via browser.

Step 2 Select the **Options** drop-down menu in the upper right corner of the page.

Step 3 Choose **Safelist**.

Step 4 From the Safelist dialog box, enter the email address or domain. You can enter multiple domains and email addresses, separated by commas.

Step 5 Click **Add to List**.

Adding Senders to Blocklists (End Users)

Messages from blocklisted senders may be rejected or quarantined, depending on the safelist/blockist action settings defined by your administrator.

Note You can add blocklist entries only using this procedure.

Step 1 Log in to the spam quarantine.

Step 2 Select the **Options** drop-down menu in the upper right corner of the page.

Step 3 Enter the domain or email address that you want to blocklist. You can enter multiple domains and email addresses, separated by commas.

Step 4 Click **Add to List**.

Backing Up and Restoring the Safelist/Blocklist

Before you upgrade your appliance or run the installation wizard, you should back up the safelist/blocklist database. Safelist/blocklist information is not included in the main XML configuration file that contains your appliance configuration settings.

Safelist/blocklist entries can also be backed up along with other data on the Security Management appliance. See Backing Up Security Management Appliance Data , on page 290.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Select **Management Appliance > System Administration > Configuration File**.

Step 3 Scroll to the **End-User Safelist/Blocklist Database (Spam Quarantine)** section.

To	Do This
Export the safelist/blocklist	Note the path and filename of the .csv file, and modify as needed. Click **Backup Now**. The appliance saves a .csv file to the /configuration directory of the appliance using the following naming convention: *slbl<serial number><timestamp>.csv*

To	Do This
Import the safelist/blocklist	**Caution** This process will overwrite all existing entries in safelists and blocklists for all users. Click **Select File to Restore**. Select the desired file from the list of files in your configuration directory. Select the safelist/blocklist backup file that you want to restore. Click **Restore**.

Troubleshooting Safelists and Blocklists

To troubleshoot issues with safelists and blocklists, you can view the log files or system alerts.

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

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

For more information about alerts, see Managing Alerts, on page 312 .

For more information about log files, see Logging, on page 347 .

Message from Safelisted Sender Was Not Delivered

Problem

Message from a safelisted sender was not delivered.

Solution

Possible causes:

- The message was dropped for malware or content violations. See Message Processing of Safelists and Blocklists , on page 180.
- If you have multiple appliances and the sender was recently added to the safelist, then safelist/blocklists might not have been synchronized at the time the message was processed. See External Spam Quarantine and Safelist/Blocklists , on page 181 .

Configuring Spam Management Features for End Users

To	See
Understand the benefits and limitations of the different authentication methods for end-user access to spam management features.	Configuring End-User Access to the Spam Quarantine , on page 189 and subsections
Allow end users to access the spam quarantine directly via browser.	Authentication Options for End Users Accessing Spam Management Features , on page 187
Send users a notification when messages addressed to them are routed to the spam quarantine. Notifications can include links for access to the spam quarantine.	Notifying End Users About Quarantined Messages , on page 191
Allow users to specify email addresses and domains of senders whom they know to be safe, and of senders whom they know to be sending spam or other unwanted mail.	Using Safelists and Blocklists to Control Email Delivery Based on Sender , on page 180

Authentication Options for End Users Accessing Spam Management Features

Note Mailbox authentication does not allow users to view messages addressed to an email alias.

For End-User Spam Quarantine Access	Do This
Directly via web browser, authentication required and Via a link in a notification, authentication required	1 In the End User Quarantine Access settings, choose **LDAP, SAML 2.0,** or **Mailbox (IMAP/POP)**. 2 In the Spam Notifications settings, deselect **Enable login without credentials for quarantine access**.
Directly via web browser, authentication required and Via a link in a notification, authentication not required	1 In the End User Quarantine Access settings, choose **LDAP, SAML 2.0,** or **Mailbox (IMAP/POP)**. 2 In the Spam Notifications settings, select **Enable login without credentials for quarantine access**.
Only via a link in a notification, authentication not required	In the End User Quarantine Access settings, choose **None** as the authentication method.

For End-User Spam Quarantine Access	Do This
No access	In the End User Quarantine Access settings, deselect **Enable End-User Quarantine Access**.

LDAP Authentication Process

1. A user enters his or her username and passphrase into the web UI login page.
2. The spam quarantine connects to the specified LDAP server either to perform an anonymous search or as an authenticated user with the specified “Server Login” DN and passphrase. For Active Directory, you will usually need to have the server connect on the “Global Catalog port” (it is in the 6000s) and you need to create a low privilege LDAP user that the spam quarantine can bind as in order to execute the search.
3. The spam quarantine then searches for the user using the specified BaseDN and Query String. When a user’s LDAP record is found, the spam quarantine then extracts the DN for that record and attempts bind to the directory using the user records’ DN and the passphrase they entered originally. If this passphrase check succeeds then the user is properly authenticated, but the spam quarantine still needs to determine which mailboxes’ contents to show for that user.
4. Messages are stored in the spam quarantine using the recipient's envelope address. After a user's passphrase is validated against LDAP, the spam quarantine then retrieves the “Primary Email Attribute” from the LDAP record to determine which envelope address they should show quarantined messages for. The “Primary Email Attribute” can contain multiple email addresses which are then used to determine what envelope addresses should be displayed from the quarantine for the authenticated user.

IMAP/POP Authentication Process

1. Depending on your mail server configuration, a user enters their username (joe) or email address (joe@example.com) and passphrase into the web UI login page. You can modify the Login Page Message to tell your users whether they should enter a full email address or just their username (see Configuring End-User Access to the Spam Quarantine , on page 189).
2. The spam quarantine connects to the IMAP or POP server and uses the entered login (either username or email address) and passphrase to try to log into the IMAP/POP server. If the passphrase is accepted then the user is considered authenticated and the spam quarantine immediately logs out of the IMAP/POP server.
3. Once the user is authenticated, the spam quarantine lists email for the user, based on the email address:
   - If you have configured the spam quarantine to specify a domain to append to bare usernames (like joe), then this domain is appended and that fully qualified email address is used to search for matching envelopes in the quarantine.
   - Otherwise, the spam quarantine uses the entered email address to search for matching envelopes.

For more information about IMAP, see the University of Washington web site:

http://www.washington.edu/imap/

SAML 2.0 Authentication Process

See section *SSO Using SAML 2.0* in the *Cisco Content Security Management Appliance Guide*

Setting Up End-User Access to the Spam Quarantine via Web Browser

DETAILED STEPS

	Command or Action	Purpose
Step 1	Understand the benefits and limitations of the different authentication methods for end-user access to spam management features.	See section *SSO Using SAML 2.0* in the *Cisco Content Security Management Appliance Guide*
Step 2	If you will authenticate end users using LDAP, configure an LDAP server profile, including the **Spam Quarantine End-User Authentication Query** settings on the **System Administration > LDAP > LDAP Server Profile** page. **Example:** `If you will authenticate end users using SAML 2.0 (SSO), configure the settings on the System Administration > SAML page.`	Integrating With LDAP, on page 229 and subsections SSO Using SAML 2.0, on page 337
Step 3	Configure end-user access to the spam quarantine.	Configuring End-User Access to the Spam Quarantine , on page 189
Step 4	Determine the URL for end-user access to the spam quarantine.	Determining the URL for End-User Access to the Spam Quarantine, on page 190

Configuring End-User Access to the Spam Quarantine

Administrative users can access the spam quarantine whether or not end-user access is enabled.

Before You Begin

See requirements in Authentication Options for End Users Accessing Spam Management Features , on page 187.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Select **Management Appliance > Centralized Services > Spam Quarantine**.

Step 3 Click **Edit Settings** .

Step 4 Scroll down to the **End-User Quarantine Access** section.

Step 5 Select **Enable End-User Quarantine Access**.

Step 6 Specify the method to use to authenticate end users when they attempt to view their quarantined messages.

Select This Option	More Information
None	—

Select This Option	More Information
Mailbox (IMAP/POP)	For sites without an LDAP directory to use for authentication, the quarantine can validate user email addresses and passphrases against a standards-based IMAP or POP server that holds their mailbox. When logging in to the spam quarantine, end users enter their full email address and mailbox passphrase. If the POP server advertises APOP support in the banner, then for security reasons (i.e., to avoid sending the passphrase in the clear) the Cisco appliance will only use APOP. If APOP is not supported for some or all users then the POP server should be reconfigured to not advertise APOP. Select SSL if you have configured your server to use it. If users enter username only, you can specify a domain to add to automatically complete the email address. Enter the domain of the envelope for users logging in to "Append Domain to Unqualified Usernames."
LDAP	Configure LDAP settings as described in the sections referenced in the Before You Begin section of this topic.
SAML 2.0	Enable single sign-on for Spam Quarantine. Before using this option, make sure that you have configured all the settings on Management Appliance > System Administration > SAML page. See section *SSO Using SAML 2.0* in the *Cisco Content Security Management Appliance Guide* .

Step 7 Specify whether or not to display message bodies before messages are released.
If this box is selected, users may not view the message body via the spam quarantine page. Instead, to view the body of a quarantined message, users must release the message and view it in their mail application (such as Microsoft Outlook). You can use this feature for policy and regulation compliance — for example, if a regulation requires that all viewed email be archived.

Step 8 Submit and commit your changes.

What to Do Next

(Optional) Customize the page that users see when they access the spam quarantine, if you have not yet done so. See setting descriptions in Enabling and Configuring the Spam Quarantine , on page 175.

Determining the URL for End-User Access to the Spam Quarantine

The URL that end users can use to directly access the spam quarantine is formed from the hostname of the machine and the settings (HTTP/S and port numbers) configured on the IP interface on which the quarantine has been enabled. For example, HTTP://mail3.example.com:82.

Which Messages an End User Sees

Generally, end users see only their own messages in the spam quarantine.

Depending on the method of access (via notification or directly via web browser) and authentication method (LDAP or IMAP/POP), users may see mail for multiple email addresses in the spam quarantine.

When LDAP authentication is used, if the Primary Email attribute has multiple values in the LDAP directory, all of those values (addresses) will be associated with the user. Therefore, quarantined messages addressed to all email addresses associated with the end user in the LDAP directory are present in the quarantine.

If the authentication method is IMAP/POP, or the user accesses the quarantine directly via a notification, then the quarantine will display only messages for that user's email address (or the address to which the notification was sent).

For information about messages that are sent to aliases of which the user is a member, see Recipient Email Mailing List Aliases and Spam Notifications , on page 192.

Notifying End Users About Quarantined Messages

You can configure the system to send a notification email to some or all users when they have spam and suspected spam messages in the spam quarantine.

By default, spam notifications list the user's quarantined messages. Notifications can also include a link that users can click in order to view their quarantined messages in the spam quarantine. These links do not expire. The user can view the quarantined messages and decide whether to have them delivered to their inbox or delete them.

Note In cluster configurations, you can choose which users receive notifications only at the machine level.

Before You Begin

- For end users to manage messages listed in notifications, they must be able to access the spam quarantine. See Configuring End-User Access to the Spam Quarantine , on page 189.
- Understand the authentication options for managing spam using notifications. See Authentication Options for End Users Accessing Spam Management Features , on page 187.
- If end users receive email at multiple aliases, see Recipient Email Mailing List Aliases and Spam Notifications , on page 192.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Select **Management Appliance > Centralized Services > Spam Quarantine**.

Step 3 Click **Edit Settings** .

Step 4 Scroll down to the **Spam Notifications** section.

Step 5 Select **Enable Spam Notification**.

Step 6 Specify options.

To customize the message body:

a) (Optional) Customize the default text and variables.
To insert a variable, place the cursor where you would like the variable inserted and then click the name of the variable in the Message Variables listing on the right. Or type in the variable.

The following message variables are expanded to the actual value for the specific end user:

- **New Message Count** (%new_message_count%) — The number of new messages since the user last logged in.
- **Total Message Count** (%total_message_count%) — The number of messages for the user in the spam quarantine.
- **Days Until Message Expires** (%days_until_expire%)
- **Quarantine URL** (%quarantine_url%) — URL to log in to the quarantine and view messages.
- **Username** (%username%)
- **New Message Table** (%new_quarantine_messages%) — A list of the user's new quarantined messages, showing sender, message subject, date, and a link to release the message. The user clicks a message subject to view the message in the spam quarantine.
- **New Message Tablewithout Subject** (%new_quarantine_messages_no_subject%) — Similar to New Message Table, but only a "View Message" link is shown in place of the subject for each message.

b) If you have enabled an authentication method in the End User Quarantine Access section on this page:

- To automatically log users in to the spam quarantine when they access it by clicking a link in a notification, select **Enable login without credentials for quarantine access**. End users can release messages simply by clicking "Release" links in the notification.
- To require users to log in to the spam quarantine when they access it by clicking a link in a notification, deselect this option. End users cannot release messages simply by clicking "Release" links in the notification.

c) Click **Preview Message** to verify that the message is as you want it to be.

Step 7 Submit and commit your changes.

What to Do Next

To ensure that end users receive these notifications, consider recommending that they add the From: address for the spam quarantine notification emails to the "whitelist" in the junk mail settings of their mail application (such as Microsoft Outlook or Mozilla Thunderbird.)

Recipient Email Mailing List Aliases and Spam Notifications

Notifications can be sent to each Envelope Recipient that has quarantined email, including mailing lists and other aliases. Each mailing list receives a single digest. If you send notifications to a mailing list, all subscribers to the list will receive the notification. Users who belong to multiple email aliases, or who belong to LDAP groups that receive notifications, or who use several email addresses, may receive multiple spam notifications. The following table shows example situations in which users may receive multiple notifications.

Table 43: Notifications per Address/Alias

User	Email Addresses	Aliases	Notifications
Sam	sam@example.com	—	1
Mary	mary@example.com	dev@example.com qa@example.com pm@example.com	4
Joe	joe@example.com, admin@example.com	hr@example.com	3

If you use LDAP authentication, you can choose not to send notifications to mailing list aliases. Or, if you choose to send spam notifications to mailing list aliases, you can prevent some occurrences of multiple notifications. .

Users who access the spam quarantine by clicking a link in a notification will not see quarantined messages for any other aliases that the end-user may have, unless the appliance is using a spam quarantine alias consolidation query for email notifications. If the notification was sent to a distribution list that is expanded after processing by the appliance, then multiple recipients may have access to the same quarantine for that list.

This means that all subscribers to a mailing list will receive the notification and can log in to the quarantine to release or delete messages. In this case, end users visiting the quarantine to view messages mentioned in a notification may find that those messages have already been deleted by other users.

Note If you do not use LDAP and you do not want your end users to receive multiple email notifications, consider disabling notifications and instead allow end users to access the quarantine directly and authenticate via LDAP or POP/IMAP.

Testing Notifications

You can test notifications by configuring a testing mail policy, and having spam quarantined for just a single user. Then, configure the spam quarantine notification settings: Select the **Enable Spam Notification** checkbox and do not select **Enable End-User Quarantine Access**. Then only the administrator configured in the **Deliver Bounced Messages To** field is notified of new spam in the quarantine.

Troubleshooting Spam Notifications

User Receives Multiple Notifications

Problem

A user receives multiple spam notifications for a single message.

Solution

Possible causes:

- The user has multiple email addresses and the spam message was sent to more than one of those addresses.
- The user is a member of one or more email aliases that received the spam message. To minimize duplications, and for more information, see Recipient Email Mailing List Aliases and Spam Notifications , on page 192.

Recipient Does Not Receive Notifications

Problem

Recipient is not receiving spam notifications.

Solution

- If notifications are being sent to the "Deliver Bounce Messages To:" address instead of to spam recipients, this means that spam notifications are enabled, but spam quarantine access is not enabled. See Authentication Options for End Users Accessing Spam Management Features , on page 187.
- Have the user check the junk mail settings of their email client.
- Check for issues with the appliance or server that you specified for **Deliver Messages Via** in Enabling and Configuring the Spam Quarantine , on page 175.

Managing Messages in the Spam Quarantine

This section explains how to work with messages in local or external spam quarantines.

Administrative users can see and manage all messages in the spam quarantine.

Accessing the Spam Quarantine (Administrative Users)

Administrative users can see and manage all messages in the spam quarantine.

Accessing the Spam Quarantine (Administrative Users)

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, choose **Quarantine** > **Spam Quarantine**. The spam quarantine opens in a separate browser window.

Step 2 Select **Email > Message Quarantine > Spam Quarantine**, then click the **Spam Quarantine** link. The spam quarantine opens in a separate browser window.

Searching for Messages in the Spam Quarantine

Step 1 Specify an envelope recipient.

Note You can enter a partial address.

Step 2 Select whether the search results should match the exact recipient you entered, or whether the results should contain, start with, or end with your entry.

Step 3 Enter a date range to search through. Click the calendar icons to select a date.

Step 4 Specify a From: address, and select whether the search results should contain, match exactly, start with, or end with the value you entered.

Step 5 Click **Search**. Messages matching your search criteria are displayed below the Search section of the page.

Searching Very Large Message Collections

If you have a very large collection of messages in the spam quarantine, and if your search terms are not narrowly defined, your query may take a very long time to return information, or it may time out.

You will be prompted to confirm whether you want to resubmit your search. Please note that having multiple large searches running simultaneously can impact performance.

Viewing Messages in the Spam Quarantine

The message listing shows messages in the spam quarantine. You can select how many messages are shown at one time. You can sort the display by clicking on the column headings. Click the same column again to reverse the sorting.

Click the subject of a message to view the message, including the body and headers. The message is displayed in the Message Details page. The first 20K of the message is displayed. If the message is longer, it is truncated at 20K and you can download the message via the link at the bottom of the message.

From the Message Details page you can delete a message (select **Delete**) or select **Release** to release the message. Releasing a message causes it to be delivered.

To view additional details about the message, click the **Message Tracking** link.

Note the following:

- **Viewing Messages with Attachments**

  When viewing a message that includes an attachment, the body of the message is displayed, followed by a list of attachments.

- **Viewing HTML Messages**

  The spam quarantine attempts to render an approximation of HTML-based messages. Images are not displayed.

- **Viewing Encoded Messages**

  Base64-encoded messages are decoded and then displayed.

Delivering Messages in the Spam Quarantine

To release a message for delivery, click the checkbox next to the message or messages that you want to release and select **Release** from the drop-down menu. Then click **Submit**.

Click the checkbox in the heading row to automatically select all messages currently displayed on the page.

Released messages proceed directly to the destination queue, skipping any further work queue processing in the email pipeline.

Deleting Messages from the Spam Quarantine

The spam quarantine can be configured to automatically delete messages after a certain amount of time. Also, the spam quarantine can be configured to automatically delete the oldest messages once the quarantine has reached its maximum size. You may also delete messages from the spam quarantine manually.

To delete specific messages, click the checkbox next to the messages that you want to delete and then select **Delete** from the drop-down menu. Then click **Submit**. Click the checkbox in the heading row to automatically select all of the messages currently displayed on the page.

To delete all messages in the spam quarantine, disable the quarantine (see About Disabling the External Spam Quarantine , on page 196) and then click the **Delete All Messages** link. The number in parenthesis at the end of the link is the number of messages in the spam quarantine.

Disk Space for the Spam Quarantine

Available disk space for the quarantine varies by appliance model. See Viewing Disk Space, Quotas and Usage , on page 334.

By default, messages in the spam quarantine are automatically deleted after a set amount of time. If the quarantine gets full, older spam is deleted. To change this setting, see Enabling and Configuring the Spam Quarantine , on page 175.

About Disabling the External Spam Quarantine

If you disable the spam quarantine:

- If messages are present in the spam quarantine when it is disabled, you can opt to delete all of the messages.
- Any mail polices set to quarantine spam or suspected spam will instead be set to deliver the message. You may need to adjust mail policies on the Email Security appliance.
- To completely disable an external spam quarantine, disable it on both the Email Security appliance and the Security Management appliance.

Disabling an external spam quarantine on the Email Security appliance only does not delete the external quarantine or its messages and data.

Troubleshooting Spam Quarantine Features

- Troubleshooting Safelists and Blocklists, on page 186
- Troubleshooting Spam Notifications , on page 193

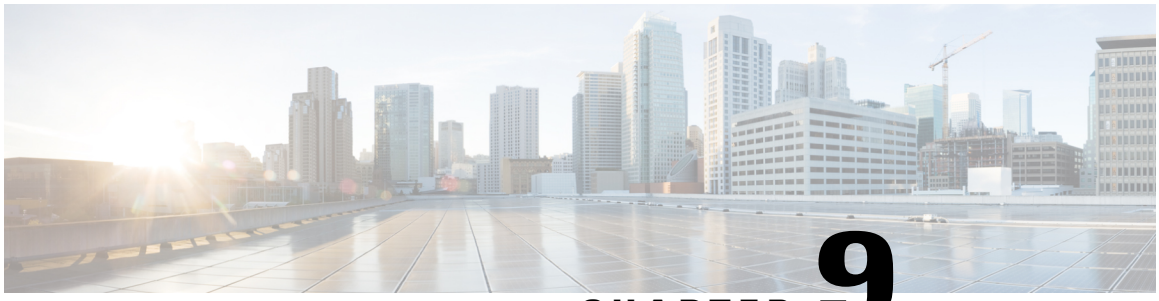

CHAPTER 9

Centralized Policy, Virus, and Outbreak Quarantines

This chapter contains the following sections:

Overview of Centralized Quarantines

Messages processed by certain filters, policies, and scanning operations on an Email Security appliance can be placed into quarantines to temporarily hold them for further action. You can centralize quarantines from multiple Email Security appliances on a Cisco Content Security Management appliance.

Benefits of centralizing quarantines include the following:

- You can manage quarantined messages from multiple Email Security appliances in one location.
- Quarantined messages are stored behind the firewall instead of in the DMZ, reducing security risk.
- Centralized quarantines can be backed up as part of the standard backup functionality on the Security Management appliance.

Anti-virus scanning, Outbreak Filters, and Advanced Malware Protection (File Analysis) each have a single dedicated quarantine. You create policy quarantines to hold messages that are caught by message filtering, content filtering, and Data Loss Prevention policies.

The Policy, Virus and Outbreak Quarantines section in the legacy web interface is labeled as Other Quarantines in the new web interface. For more information, see Viewing Messages in Quarantines , on page 215.

For additional information about quarantines, see the documentation for your Email Security appliance.

Quarantine Types

Quarantine Type	Quarantine Name	Created by the System by Default?	Description	More Information
Advanced Malware Protection	File Analysis	Yes	Holds messages that are sent for file analysis, until a verdict is returned.	• Managing Policy, Virus, and Outbreak Quarantines, on page 206 • Working with Messages in Policy, Virus, or Outbreak Quarantines, on page 215
Virus	Virus	Yes	Holds messages that may be transmitting malware, as determined by the anti-virus engine.	
Outbreak	Outbreak	Yes	Holds messages caught by Outbreak Filters as potentially being spam or malware.	
Policy	Policy	Yes	Holds messages caught by message filters, content filters, and DLP message actions. A default Policy quarantine has been created for you.	
	Unclassified	Yes	Holds messages only if a quarantine that is specified in a message filter, content filter, or DLP message action has been deleted. You cannot assign this quarantine to any filter or message action.	
	(Policy quarantines that you create)	No	Policy quarantines that you create for use in message filters, content filters, and DLP message actions.	

Quarantine Type	Quarantine Name	Created by the System by Default?	Description	More Information
Spam	Spam	Yes	Holds spam or suspected spam messages for the message's recipient or an administrator to review. The spam quarantine is not included in the group of policy, virus, and outbreak quarantines and is managed separately from all other quarantines.	Spam Quarantine, on page 173

Centralizing Policy, Virus, and Outbreak Quarantines

DETAILED STEPS

	Command or Action	Purpose
Step 1	If your Email Security appliance is in your DMZ and your Security Management appliance is behind your firewall, open a port in the firewall to allow the appliances to exchange centralized policy, virus, and outbreak quarantine data.	Firewall Information, on page 401
Step 2	On the Security Management appliance, enable the feature.	Enabling Centralized Policy, Virus, and Outbreak Quarantines on the Security Management Appliance , on page 201
Step 3	On the Security Management appliance, allocate disk space for non-spam quarantines.	Managing Disk Space , on page 333
Step 4	(Optional) • Create centralized policy quarantines on the Security Management appliance with desired settings. • Configure settings for the centralized virus and outbreak quarantines, and for the default policy quarantines. If you configure these settings before migration, you can refer to the existing settings on your Email Security appliances. You can also create required quarantines while configuring custom migration, or quarantines will be created for you during automatic	• Configuring Policy, Virus, and Outbreak Quarantines , on page 208 • Checking the Settings of System-Created Quarantines , on page 208.

	Command or Action	Purpose
	migration. All quarantines created during migration have default settings. Local quarantine settings are not retained in the centralized quarantine, even if the quarantine name is the same.	
Step 5	On the Security Management appliance, add Email Security appliances to manage, or select the Policy, Virus and Outbreak Quarantines option from the centralized services of an already-added appliance. • If your Email Security appliances are clustered, all appliances that belong to a particular level (machine, group, or cluster) must be added to the Security Management appliance before you enable centralized Policy, Virus and Outbreak Quarantines on any Email Security appliance in the cluster.	Adding the Centralized Policy, Virus, and Outbreak Quarantine Service to Each Managed Email Security Appliance, on page 202
Step 6	Commit your changes.	
Step 7	On the Security Management appliance, configure migration of existing policy quarantines from Email Security appliances.	Configuring Migration of Policy, Virus, and Outbreak Quarantines , on page 203
Step 8	On an Email Security appliance, enable the centralized policy, virus, and outbreak quarantines feature. • **Important** If you have policy, virus, and outbreak quarantines configured on an Email Security appliance, migration of quarantines and all their messages begins as soon as you commit this change.	See the "Centralizing Services on a Cisco Content Security Management appliance" chapter in the documentation for your Email Security appliance, specifically the following sections: • "About Migration of Policy, Virus, and Outbreak Quarantines" • "Centralizing Policy, Virus, and Outbreak Quarantines"
Step 9	Migrate additional Email Security appliances. • Only one migration process can be in progress at any time. Do not enable centralized policy, virus, and outbreak quarantines on another Email Security appliance until the previous migration is complete.	
Step 10	Edit centralized quarantine settings as needed. • Quarantines created during migration are created with default settings, not the settings in the originating local quarantines, even if the centralized and local quarantine names are the same.	Configuring Policy, Virus, and Outbreak Quarantines , on page 208
Step 11	If message filters, content filters, and DLP message actions could not be automatically updated with the names of centralized quarantines, manually update those configurations on your Email Security appliances.	See the documentation for message filters, content filters, and DLP Message Actions in the online help or user guide for your Email Security appliance.

	Command or Action	Purpose
	• In cluster configurations, filters and message actions can be automatically updated on a particular level only if filters and message actions are defined at that level.	
Step 12	(Recommended) Specify an Email Security appliance to process released messages if the originating appliance is not available.	Designating an Alternate Appliance to Process Released Messages, on page 205
Step 13	If you delegate administration to custom user roles, you may need to configure access in a certain way.	Configuring Centralized Quarantine Access for Custom User Roles , on page 205

Enabling Centralized Policy, Virus, and Outbreak Quarantines on the Security Management Appliance

Before You Begin

Complete any steps preceding this procedure in the table in Centralizing Policy, Virus, and Outbreak Quarantines , on page 199.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Policy, Virus, and Outbreak Quarantines**.

Step 3 Click **Enable**.

Step 4 Specify the interface and port for communication with Email Security appliances:

- Accept the default selections unless you have reason to change them.
- If your Email Security appliances are not on the same network as your Security Management appliance, then you must use the Management interface.
- Use the same port that you opened in the firewall.

Step 5 Click **Submit**.

What to Do Next

Return to the next step in the table in Centralizing Policy, Virus, and Outbreak Quarantines , on page 199.

Adding the Centralized Policy, Virus, and Outbreak Quarantine Service to Each Managed Email Security Appliance

To see an consolidated view of all quarantines on all Email Security appliances, consider adding all Email Security appliances before centralizing any quarantines.

Before You Begin

Make sure you have completed all procedures to this point in the table in Centralizing Policy, Virus, and Outbreak Quarantines , on page 199.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Security Appliances**.

Step 3 If you have already added the Email Security appliance to the list on this page:

a) Click the name of an Email Security appliance.

b) Select the **Policy, Virus, and Outbreak Quarantines** service.

Step 4 If you have not yet added the Email Security appliance:

a) Click Add Email Appliance.

b) In the Appliance Name and IP Address text fields, enter the appliance name and the IP address for the Management interface of the appliance you are adding.

Note If you enter a DNS name in the IP Address text field, it will be immediately resolved to an IP address when you click **Submit**.

c) The Policy, Virus and Outbreak Quarantines service is pre-selected.

d) Click **Establish Connection**.

e) Enter the user name and password for an administrator account on the appliance to be managed, and then click **Establish Connection**.

Note You enter the login credentials to pass a public SSH key for file transfers from the Security Management appliance to the remote appliance. The login credentials are not stored on the Security Management appliance.

f) Wait for the Success message to appear above the table on the page.

Step 5 Click **Submit**.

Step 6 Repeat this procedure for each Email Security appliance for which you want to enable Centralized Policy, Virus, and Outbreak Quarantines.
For example, add the other appliances in the cluster.

Step 7 Commit your changes.

What to Do Next

Return to the next step in the table in Centralizing Policy, Virus, and Outbreak Quarantines , on page 199.

Configuring Migration of Policy, Virus, and Outbreak Quarantines

Before You Begin

- Make sure that you have completed all procedures to this point in the table in Centralizing Policy, Virus, and Outbreak Quarantines , on page 199
- For caveats and information about the migration process, see the "About Migration of Policy, Virus, and Outbreak Quarantines" section in the "Centralizing Services on a Cisco Content Security Management appliance" chapter in the documentation for your Email Security appliance.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Policy, Virus, and Outbreak Quarantines**.

Step 3 Click **Launch Migration Wizard**.

Step 4 Choose a migration method:

If	Choose	Additional Information
• You want to migrate all existing policy quarantines from all associated Email Security appliances, and • Policy quarantines with the same names have identical settings on all Email Security appliances, and • You want to merge all policy quarantines with the same name on all Email Security appliances into a single centralized policy quarantine having that name.	Automatic	All centralized policy quarantines that are created using this process are automatically configured with default settings, regardless of the settings in the quarantines with the same names on the Email Security appliance. You must update those settings after migration.

If	Choose	Additional Information
• Policy quarantines with the same names have different settings on different Email Security appliances and you want to maintain the differences, or • You want to migrate some local quarantines and delete all others, or • You want to migrate local quarantines to centralized quarantines with different names or • You want to merge local quarantines with different names into a single centralized quarantine.	Custom	Any centralized policy quarantines that you create during migration, instead of before migration, will be configured with the default settings for new quarantines. You should update those settings after migration.

Step 5 Click **Next**.

Step 6 If you selected **Automatic**:
Verify that the policy quarantines to be migrated and other information on this page match your expectations.

Virus, Outbreak, and File Analysis quarantines will also be migrated.

Step 7 If you selected **Custom**:

- To select whether to show quarantines from all Email Security appliances or just one., choose an option from the **Show Quarantines from**: list.
- Select which local policy quarantines move to each centralized policy quarantine.
- Create additional centralized policy quarantines as needed. These will have default settings.
- Quarantine names are case-sensitive.
- Any quarantines remaining in the table on the left will not be migrated and will be deleted from the Email Security appliance upon migration.
- You can change the quarantine mapping by selecting a quarantine from the table on the right and clicking **Remove from Centralized Quarantine**.

Step 8 Click **Next** as needed.

Step 9 Submit and commit your changes.

What to Do Next

Return to the next step in the table in Centralizing Policy, Virus, and Outbreak Quarantines , on page 199.

Designating an Alternate Appliance to Process Released Messages

Normally, when a message is released from a centralized quarantine, the Security Management appliance returns it for processing to the Email Security appliance that originally sent it to the centralized quarantine.

If the Email Security appliance that originated a message is not available, a different Email Security appliance can process and deliver released messages. You designate the appliance for this purpose.

Before You Begin

- Verify that the alternate appliance can process and deliver released messages as expected. For example, configurations for encryption and antivirus rescanning should match the same configurations on your primary appliances.
- The alternate appliance must be fully configured for centralized policy, virus, and outbreak quarantines. Complete the steps in the table in Centralizing Policy, Virus, and Outbreak Quarantines , on page 199 for that appliance.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > Security Appliances**.

Step 3 Click the **Specify Alternate Release Appliance** button.

Step 4 Choose an Email Security appliance.

Step 5 Submit and commit your changes.

What to Do Next

Related Topics

Releasing Messages When an Email Security Appliance Is Unavailable , on page 206

Configuring Centralized Quarantine Access for Custom User Roles

In order to allow administrators with custom user roles to specify centralized policy quarantines in message and content filters and in DLP message actions on the Email Security appliance, you must grant those users access to the relevant policy quarantines on the Security Management appliance, and the custom user role names that you create on the Security Management appliance must match those on the Email Security appliance.

Related Topics

- Creating Custom Email User Roles , on page 258

Disabling Centralized Policy, Virus, and Outbreak Quarantines

Generally, if you need to disable these centralized quarantines, you should do so on the Email Security appliance.

For information about disabling centralized policy, virus, and outbreak quarantines, including a list of impacts of doing so, see the online help or documentation for your Email Security appliance.

Releasing Messages When an Email Security Appliance Is Unavailable

Normally, when a message is released from a centralized quarantine, the Security Management appliance returns it for processing to the Email Security appliance that originally sent it to the centralized quarantine.

If the Email Security appliance that originated a message is not available, a different Email Security appliance can process and deliver released messages. You should designate an alternate release appliance for this purpose.

If the alternate appliance is unavailable, you can specify a different Email Security appliance as the alternate release appliance and that appliance will process and deliver queued messages.

After repeated unsuccessful attempts to reach an Email Security appliance. you will receive an alert.

Related Topics

- Designating an Alternate Appliance to Process Released Messages, on page 205

Managing Policy, Virus, and Outbreak Quarantines

Disk Space Allocation for Policy, Virus, and Outbreak Quarantines

For information about allocating disk space, see Managing Disk Space , on page 333.

Messages in multiple quarantines consume the same amount of disk space as a message in a single quarantine.

If Outbreak Filters and Centralized Quarantines are both enabled:

- All disk space on the Email Security appliance that would have been allocated to local policy, virus, and outbreak quarantines is used instead to hold copies of messages in the Outbreak quarantine, in order to scan those messages each time outbreak rules are updated.
- The disk space on the Security Management appliance for messages in the Outbreak quarantine from a particular managed Email Security appliance may be limited by the amount of available disk space for quarantined messages on that Email Security appliance.
- For more information about this situation, see Retention Time for Messages in Quarantines , on page 206

Retention Time for Messages in Quarantines

Messages are automatically removed from the quarantine under the following circumstances:

- Normal Expiration—the configured retention time is met for a message in the quarantine. You specify a retention time for messages in each quarantine. Each message has its own specific expiration time, displayed in the quarantine listing. Messages are stored for the amount of time specified unless another circumstance described in this topic occurs.

  Note The normal retention time for messages in the Outbreak Filters quarantine is configured in the Outbreak Filters section of each mail policy, not in the outbreak quarantine.

- Early Expiration—messages are forced from quarantines before the configured retention time is reached. This can happen when:

  - The size limit for all quarantines, as defined in Disk Space Allocation for Policy, Virus, and Outbreak Quarantines , on page 206, is reached.

    If the size limit is reached, the oldest messages, regardless of quarantine, are processed and the default action is performed for each message, until the size of all quarantines is again less than the size limit. The policy is First In First Out (FIFO). Messages in multiple quarantines will be expired based on their latest expiration time.

    (Optional) You can configure individual quarantines to be exempt from release or deletion because of insufficient disk space. If you configure all quarantines to be exempt and the disk space reaches capacity messages will be held on the Email Security appliance until space is available on the Security Management appliance.

    Because the Security Management appliance does not scan messages, a copy of each message in the centralized outbreak quarantine is stored on the Email Security appliance that originally processed the message. This allows the Email Security appliance to rescan quarantined messages each time outbreak filter rules are updated, and tell the Security Management appliance to release messages that are no longer deemed a threat. Both copies of the outbreak quarantine should hold the same set of messages at all times. Therefore, in the rare situation when disk space on the Email Security appliance becomes full, then the copies of messages in the Outbreak quarantine on both appliances will expire early, even if the centralized quarantine still has space.

    You will receive alerts at disk-space milestones. See Alerts About Quarantine Disk-Space Usage , on page 213.

- You delete a quarantine that still holds messages.

When a message is automatically removed from a quarantine, the default action is performed on that message. See Default Actions for Automatically Processed Quarantined Messages , on page 208.

Note In addition to the above scenarios, messages can be automatically removed from quarantine based on the result of scanning operations (outbreak filters or file analysis.)

Effects of Time Adjustments on Retention Time

- Daylight savings time and appliance time zone changes do not affect the retention period.
- If you change the retention time of a quarantine, only new messages will have the new expiration time.
- If the system clock is changed, messages that should have expired in the past will expire at the next most appropriate time.

- System clock changes do not apply to messages that are in the process of being expired.

Default Actions for Automatically Processed Quarantined Messages

The default action is performed on messages in a policy, virus, or outbreak quarantine when any situation described in Retention Time for Messages in Quarantines , on page 206, occurs.

There are two primary default actions:

- Delete—The message is deleted.
- Release—The message is released for delivery.

Upon release, messages may be rescanned for threats. For more information, see About Rescanning of Quarantined Messages , on page 220.

In addition, messages released before their expected retention time has passed can have additional operations performed on them, such as adding an X-Header. For more information, see Configuring Policy, Virus, and Outbreak Quarantines , on page 208.

Messages released from a centralized quarantine are returned to the originating Email Security appliance for processing.

Checking the Settings of System-Created Quarantines

Before you use quarantines, customize the settings of the default quarantines, including the Unclassified quarantine.

Configuring Policy, Virus, and Outbreak Quarantines

Before You Begin

- If you are editing an existing quarantine, see About Editing Policy, Virus, and Outbreak Quarantine Settings , on page 210.
- Understand how messages in quarantines are automatically managed, including retention times and default actions. See Retention Time for Messages in Quarantines , on page 206, and Default Actions for Automatically Processed Quarantined Messages , on page 208.
- Determine which users you want to have access to each quarantine, and create users and custom user roles accordingly. For details, see Which User Groups Can Access Policy, Virus, and Outbreak Quarantines , on page 214.

Step 1 [New Web Interface Only] Choose **Quarantine** > **Other Quarantine** > **View** > +.

Step 2 Choose **Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines**.

Step 3 Do one of the following:

- Click **Add Policy Quarantine**.
- Click a quarantine to edit.

Step 4 Enter information.
Keep the following in mind:

- Changing the retention time of the File Analysis quarantine from the default of one hour is not recommended.
- If you do *not* want messages in this quarantine to be processed before the end of the Retention Period you specify, even when quarantine disk space is full, deselect **Free up space by applying default action on messages upon space overflow**.

  Do not select this option for all quarantines. The system must be able to make space by deleting messages from at least one quarantine.
- If you select **Release** as the default action, you can specify additional actions to apply to messages that are released before their retention period has passed:

Option	Information
Modify Subject	Type the text to add and specify whether to add it to the beginning or the end of the original message subject. For example, you might want to warn the recipient that the message may contain inappropriate content. **Note** In order for a subject with non-ASCII characters to display correctly it must be represented according to RFC 2047.
Add X-Header	An X-Header can provide a record of actions taken on a message. This can be helpful for example when handling inquiries about why a particular message was delivered. Enter a name and value. Example: Name = Inappropriate-release-early Value = True
Strip Attachments	Stripping attachments protects against viruses that may be in such files.

Step 5 Specify the users who can access this quarantine:

User	Information
Local Users	The list of local users includes only users with roles that can access quarantines. The list excludes users with Administrator privileges, because all Administrators have full access to quarantines.
Externally Authenticated Users	You must have configured external authentication.
Custom User Roles	You see this option only if you have created at least one custom user role with quarantine access.

Step 6 Submit and commit your changes.

What to Do Next

See Message Filters Page, on page 143 and Content Filters Page, on page 144

- If you have not yet migrated quarantines from the Email Security appliance:
  You will assign these quarantines to message and content filters and DLP message actions as part of the migration process.
- If you have already migrated to centralized quarantines:

  Make sure your Email Security appliance has message and content filters and DLP message actions that will move messages to the quarantine. See the user guide or online help for the Email Security appliance.

About Editing Policy, Virus, and Outbreak Quarantine Settings

Note

- You cannot rename a quarantine.
- See also Retention Time for Messages in Quarantines , on page 206.

To change quarantine settings, choose Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines from the Appliance Configuarion page, and then click the name of a quarantine.

Determining the Filters and Message Actions to Which a Policy Quarantine Is Assigned

You can view the message filters, content filters, Data Loss Prevention (DLP) message actions, and DMARC verification profiles that are associated with a policy quarantine, and the Email Security appliance on which each is configured.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines**.

Step 3 Click the name of the policy quarantine to check.

Step 4 Scroll to the bottom of the page and view the **Associated Message Filters/Content Filters/DLP Message Actions**.

About Deleting Policy Quarantines

- Before you delete a policy quarantine, see if it is associated with any active filters or message actions. See Determining the Filters and Message Actions to Which a Policy Quarantine Is Assigned , on page 210.
- You can delete a policy quarantine even if it is assigned to a filter or message action.
- If you delete a quarantine that is not empty, the default action defined in the quarantine will be applied to all messages, even if you have selected the option not to delete messages if the disk is full. See Default Actions for Automatically Processed Quarantined Messages , on page 208.
- After you delete the quarantine associated with a filter or message action, any messages subsequently quarantined by that filter or message action will be sent to the Unclassified quarantine. You should customize the default settings of the Unclassified quarantine before you delete quarantines.
- You cannot delete the Unclassified quarantine.

Monitoring Quarantine Status, Capacity, and Activity

To View	Do This
Total space allocated for all non-spam quarantines	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear ⚙ icon to load the legacy web interface. Choose **Management Appliance > Centralized Services > Policy, Virus, and Outbreak Quarantines** and look in the first section on the page. To change allocations, see Managing Disk Space , on page 333 .
Currently available space for all non-spam quarantines	[New Web Interface Only] Choose **Quarnatine** > **Other Quarantine**. or Choose **Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines** and look just below the table. The available space for policy, virus and outbreak quarantine is displayed above the table in the Quarantines section
Total amount of space currently used by all quarantines	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear ⚙ icon to load the legacy web interface. Choose **Management Appliance > Centralized Services > System Status**.

To View	Do This
Amount of space currently used by each quarantine	[New Web Interface Only] Choose **Quarantines Quarantine** > **Other Quarantine** > **View**. The table displays the amount of space currently used by each quarantine. or Choose **Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines**, click the quarantine name, and look for this information in the table row directly below the quarantine name.
Total number of messages currently in all quarantines	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear ⚙ icon to load the legacy web interface. Choose **Management Appliance > Centralized Services > System Status**.
Number of messages currently in each quarantine	[New Web Interface Only] Choose **Quarantines Quarantine** > **Other Quarantine** > **View**. The table displays the total number of messages currently available for each quarantine. or Choose **Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines** and look at the table row for the quarantine.
Total CPU usage by all quarantines	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear ⚙ icon to load the legacy web interface. Choose **Management Appliance > Centralized Services > System Status** and look in the System Information section.
Date and time when the last message entered each quarantine (excluding moves between policy quarantines)	[New Web Interface Only] Choose **Quarantines Quarantine** > **Other Quarantine** > **View**. The table displays the date and time when the last message was quarantined. or Choose **Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines** and look at the table row for the quarantine.

To View	Do This
Date a policy quarantine was created	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear ⚙ icon to load the legacy web interface. Choose **Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines**, click the quarantine name, and look for this information in the table row directly below the quarantine name. Creation date and creator name are not available for system-created quarantines.
Name of policy quarantine creator	
Filters and message actions associated with a policy quarantine	See Determining the Filters and Message Actions to Which a Policy Quarantine Is Assigned , on page 210.

Alerts About Quarantine Disk-Space Usage

An alert is sent whenever the total size of the policy, virus, and outbreak quarantine reaches or passes 75 percent, 85 percent, and 95 percent of its capacity. The check is performed when a message is placed in the quarantine. For example, if adding a message to a quarantine increases the size to or past 75 percent of the total capacity, an alert is sent.

For more information about Alerts, see Managing Alerts, on page 312.

Policy Quarantines and Logging

AsyncOS individually logs all messages that are quarantined:

Info: MID 482 quarantined to "Policy" (message filter:policy_violation)

The message filter or Outbreak Filters feature rule that caused the message to be quarantined is placed in parentheses. A separate log entry is generated for each quarantine in which the message is placed.

AsyncOS also individually logs messages that are removed from quarantine:

Info: MID 483 released from quarantine "Policy" (queue full)

Info: MID 484 deleted from quarantine "Anti-Virus" (expired)

The system individually logs messages after they are removed from all quarantines and either permanently deleted or scheduled for delivery, for example

Info: MID 483 released from all quarantines

Info: MID 484 deleted from all quarantines

When a message is re-injected, the system creates a new Message object with a new Message ID (MID). This is logged using an existing log message with a new MID “byline”, for example:

Info: MID 483 rewritten to 513 by Policy Quarantine

About Distributing Message Processing Tasks to Other Users

You can distribute message review and processing tasks to other administrative users. For example:

- The Human Resources team can review and manage the Policy Quarantine.
- The Legal team can manage the Confidential Material Quarantine.

You assign access privileges to these users when you specify settings for a quarantine. In order to add users to quarantines, the users must already exist.

Each user may have access to all, some, or none of the quarantines. A user who is not authorized to view a quarantine will not see any indication of its existence anywhere in the GUI or CLI listings of quarantines.

Which User Groups Can Access Policy, Virus, and Outbreak Quarantines

When you allow administrative users to access a quarantine, the actions that they can perform depend on their user group:

- Users in the Administrators or Email Administrators groups can create, configure, delete, and centralize quarantines and can manage quarantined messages.
- Users in the Operators, Guests, Read-Only Operators, and Help Desk Users groups, as well as custom user roles with quarantine management privileges, can search for, view, and process messages in a quarantine, but cannot change the quarantine's settings, create, delete, or centralize quarantines. You specify in each quarantine which of these users have access to that quarantine.
- Users in the Technicians group cannot access quarantines.

Access privileges for related features, such as Message Tracking and Data Loss Prevention, also affect the options and information that an administrative user sees on Quarantine pages. For example, if a user does not have access to Message Tracking, that user will not see message trackinginformation for quarantined messages.

Note: To allow custom user roles configured on the Security Management appliance to specify policy quarantines in filters and DLP message actions, see Configuring Centralized Quarantine Access for Custom User Roles , on page 205.

End users do not have see or have access to policy, virus, and outbreak quarantines.

Working with Messages in Policy, Virus, or Outbreak Quarantines

Viewing Messages in Quarantines

To	Do This
View all messages in a quarantine	[New Web Interface Only] Choose **Quarantine** > **Other Quarantine** > **View**. or Choose **Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines**. In the row for the relevant quarantine, click the blue number in the **Messages** column of the table.
View messages in the Outbreak quarantine	[New Web Interface] Choose **Quarantine** > **Other Quarantine** > **View**. or Choose **Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines**. In the row for the relevant quarantine, click the blue number in the **Messages** column of the table. See Manage by Rule Summary Link, on page 221 or Rule Summary View, on page 221 [New Web Interface Only].
Navigate through the list of messages in a quarantine	Click Previous, Next, a page number, or double-arrow link. The double arrows take you to the first (<<) or last (>>) page in the listing. [New Web Interface Only] Scroll down on the table to display details of all the new messages.
Sort the list of messages in a quarantine	Click a column heading (except columns that could include multiple items or the “In quarantines” column).
Resize table columns	Drag the divider between column headings.
Customize table columns	Click  and select the columns to display, and click Close
View the content that caused the message to be quarantined	See Viewing Matched Content , on page 219.

Quarantined Messages and International Character Sets

For messages with subjects that contain characters from international character sets (double-byte, variable length, and non-ASCII encoded), the Policy Quarantine pages display subject lines in non-ASCII characters in their decoded form.

Searching for Messages in Policy, Virus, and Outbreak Quarantines

Note

- Users can find and see only the messages in quarantines to which they have access.
- Searches in Policy, Virus, and Outbreak quarantines do not find messages in the spam quarantine.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, choose **Quarantine** > **Other Quarantine** > **Search**.

Step 2 [New Web Interface Only] Click on the blue number link of the corresponding quarantine.

Tip [New Web Interface Only] For the Outbreak Quarantine, you can also find all messages quarantined by each outbreak rule: Click the **Rule Summary** tab in the Outbreak quarantine, and then click the relevant rule.

Step 3 Choose **Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines**.

Step 4 Click the **Search Across Quarantines** button.

Tip For the Outbreak Quarantine, you can also find all messages quarantined by each outbreak rule: Click the **Manage by Rule Summary** link in the Outbreak table row, and then click the relevant rule.

Step 5 (Optional) Enter other search criteria.

- For Envelope Sender and Envelope Recipient: You can enter any character(s). No validation of your entry is performed.
- Search results include only messages that match *all* of the criteria you specify. For example, if you specify an Envelope Recipient and a Subject, only messages that match the terms specified in both the Envelope Recipient *and* the Subject are returned.

What to Do Next

You can use the search results in the same way that you use the quarantine listings. For more information, see Manually Processing Messages in a Quarantine, on page 217.

For information on modifying your search criteria, see Modifying Search Criteria, on page 216.

Modifying Search Criteria

You can modify the search criteria to a custom time range or a different quarantine.

To modify the search criteria, click **Modify**.

Manually Processing Messages in a Quarantine

Manually processing messages means to manually select a Message Action for the message from the Message Actions page.

You can perform the following actions on messages:

- Delete
- Release
- Delay Scheduled Exit from quarantine
- Send a Copy of messages to email addresses that you specify
- Move a message from one quarantine to another

Generally, you can perform actions on messages in the lists that are displayed when you do the following. However, not all actions are available in all situations.

- From the list of quarantines on the **Email > Message Quarantine > Policy, Virus, and Outbreak Quarantines** page or [New Web Interface Only] **Quarantine** > **Other Quarantine** > **View** page, click the number of messages in a quarantine.
- Click on the check box of the quarantine message and select the required actions.

You can perform these actions on multiple messages at one time by:

- Choosing an option from the pick list at the top of the list of messages.
- Selecting the check box beside each message listed on a page.
- Selecting the check box in the table heading at the top of a list of messages. This applies the action to all messages visible on the screen. Messages on other pages are not affected.

Additional options are available for messages in the outbreak quarantine. See information about the Rule Summary view in the chapter on Outbreak Filters in the online help or user guide for the *AsyncOS for Email Security Appliances*.

Sending a Copy of the Message

Only users who belong to the Administrators group may send copies of a message.

To send a copy of the message, enter an email address in the Send Copy To: field and click **Submit**. Sending a copy of a message does not cause any other action to be performed on the message.

About Moving Messages Between Policy Quarantines

You can manually move messages from one policy quarantine to another on a single appliance.

When you move a message to a different quarantine:

- The expiration time is unchanged. The message keeps the retention time of the original quarantine.

- The reason the message was quarantined, including the matched content and other relevant details, does not change.
- If a message is in multiple quarantines and you move the message to a destination that already holds a copy of that message, the expiration time and reason for quarantine of the moved copy of the message overwrite those of the copy of the message that was originally in the destination quarantine.

Messages in Multiple Quarantines

If a message is present in one or more other quarantines, the "In other quarantines" column in the quarantine message list will show "Yes," regardless of whether you have permissions to access those other quarantines.

A message in multiple quarantines:

- Is not delivered unless it has been released from all of the quarantines in which it resides. If it is deleted from any quarantine, it will never be delivered.
- Is not deleted from any quarantine until it has been deleted or released from all quarantines in which it resides.

Because a user wanting to release a message may not have access to all of the quarantines in which it resides, the following rules apply:

- A message is not released from any quarantine until it has been released from all of the quarantines in which it resides.
- If a message is marked as Deleted in any quarantine, it cannot be delivered from any other quarantine in which it resides. (It can still be released.)

If a message is queued in multiple quarantines and a user does not have access to one or more of the other quarantines:

- The user will be informed whether the message is present in each of the quarantines to which the user has access.
- The GUI shows only the scheduled exit time from the quarantines to which the user has access. (For a given message, there is a separate exit time for each quarantine.)
- The user will not be told the names of the other quarantine(s) holding the message.
- The user will not see matched content that caused the message to be placed into quarantines that the user does not have access to.
- Releasing a message affects only the queues to which the user has access.
- If the message is also queued in other quarantines not accessible to the user, the message will remain in quarantine, unchanged, until acted upon by users who have the required access to the remaining quarantines (or until the message is released "normally" via early or normal expiration).

Message Details and Viewing Message Content

Click on the subject line of a message to view that message's content and to access the Quarantined Message page.

The Quarantined Message page has two sections: Quarantine Details and Message Details.

From the Quarantined Message page, you can read the message, select a Message Action, or send a copy of the message. You can also see if a message will be encrypted upon release from the quarantine due to the Encrypt on Delivery filter action.

The Message Details section displays the message body, message headers, and attachments. Only the first 100 K of the message body is displayed. If the message is longer, the first 100 K is shown, followed by an ellipsis (...). The actual message is not truncated. This is for display purposes only. You can download the message body by clicking [message body] in the Message Parts section at the bottom of Message Details. You can also download any of the message's attachments by clicking the attachment's filename.

Note The maximum limit for the attachment downloads on the Message Details page is restricted to 25 MB.

If you view a message that contains a virus and you have desktop anti-virus software installed on your computer, your anti-virus software may complain that it has found a virus. This is not a threat to your computer and can be safely ignored.

To view additional details about the message, click the **Message Tracking** link.

Note For the special Outbreak quarantine, additional functionality is available. See The Outbreak Quarantine, on page 221.

Viewing Matched Content

When you configure a quarantine action for messages that match Attachment Content conditions, Message Body or Attachment conditions, Message body conditions, or the Attachment content conditions, you can view the matched content in the quarantined message. When you display the message body, the matched content is highlighted in yellow, except for DLP policy violation matches. You can also use the $MatchedContent action variable to include the matched content from message or content filter matches in the message subject.

If the attachment contains the matched content, the attachment's contents are displayed, as well as the reason it was quarantined, whether it was due to a DLP policy violation, content filter condition, message filter condition, or Image Analysis verdict.

When you view messages in the local quarantine that have triggered message or content filter rules, the GUI may display content that did not actually trigger the filter action (along with content that triggered the filter action). The GUI display should be used as a guideline for locating content matches, but does not necessarily reflect an exact list of content matches. This occurs because the GUI uses less strict content matching logic

than is used in the filters. This issue applies only to the highlighting in the message body. The table that lists the matched strings in each part of the message, along with the associated filter rule, is correct.

Figure 4: Matched Content Viewed in the Policy Quarantine

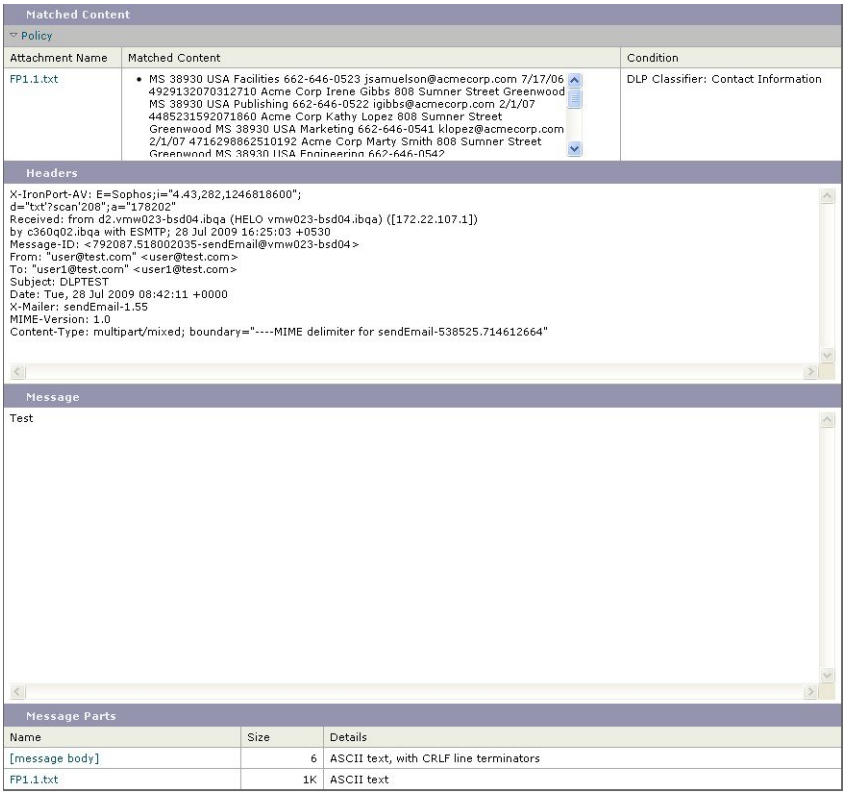

Downloading Attachments

You can download a message attachment by clicking the attachment's file name in the Message Parts or Matched Content section. AsyncOS displays a warning that attachments from unknown sources may contain viruses and asks you if you want to continue. Download attachments that may contain viruses at your own risk. You can also download the message body by clicking [message body] in the Message Parts section.

About Rescanning of Quarantined Messages

When a message is released from all queues in which is has been quarantined, the following rescanning occurs, depending on the features enabled for the appliance and for the mail policy that originally quarantined the message:

- Messages released from Policy and Virus quarantines are rescanned by the anti-virus engine.
- Messages released from the Outbreak quarantine are rescanned by the anti-spam and anti-virus engines.
- Messages released from the File Analysis quarantine are rescanned for threats.
- Messages with attachments are rescanned by the file reputation service upon release from Policy, Virus, and Outbreak quarantines.

Upon rescanning, if the verdict produced matches the verdict produced the previous time the message was processed, the message is not re-quarantined. Conversely, if the verdict is different, the message could be sent to another quarantine.

The rationale is to prevent messages from looping back to the quarantine indefinitely. For example, suppose a message is encrypted and therefore sent to the Virus quarantine. If an administrator releases the message, the anti-virus engine will still not be able to decrypt it; however, the message should not be re-quarantined or a loop will be created and the message will never be released from the quarantine. Since the two verdicts are the same, the system bypasses the Virus quarantine the second time.

The Outbreak Quarantine

The Outbreak quarantine is present when a valid Outbreak Filters feature license key has been entered. The Outbreak Filters feature sends messages to the Outbreak quarantine, depending on the threshold set. For more information, see the Outbreak Filters chapter in the online help or user guide for the Email Security appliance.

The Outbreak quarantine functions just like other quarantines—you can search for messages, release or delete messages, and so on.

The Outbreak quarantines has the following views:

The Outbreak quarantine has some additional features not available in other quarantines: the **Rule Summary** view, the **Send to Cisco** feature when viewing message details, and the option to sort messages in search results by the **Scheduled Exit** time.

If the license for the Outbreak Filters feature expires, you will be unable to add more messages to the Outbreak quarantine. Once the messages currently in the quarantine have expired and the Outbreak quarantine becomes empty, it is no longer shown in the Quarantines listing in the GUI.

Rescanning Messages in an Outbreak Quarantine

Messages placed in the Outbreak quarantine are automatically released if newly published rules deem the quarantined message no longer a threat.

If anti-spam and anti-virus are enabled on the appliance, the scanning engines scan every message released from the Outbreak quarantine based on the mail flow policy that applies to the message.

Rule Summary View

The Rule Summary view is available only in the new web interface.

In the Outbreak quarantine, click the **Rule Summary** tab to see the listing of the contents of Outbreak quarantine, grouped by Rule ID.

You can perform message actions (Release and Delete) on all of the messages in the quarantine based on which outbreak rule caused the message to be quarantined. This is ideal for clearing out large numbers of messages from the Outbreak quarantine. For more information, see Outbreak Quarantine and the Manage by Rule Summary View section of chapter “Outbreak Filters” in the online help or user guide of the *AsyncOS for the Email Security Appliance*.

Manage by Rule Summary Link

Click the Manage by Rule Summary link next to the Outbreak quarantine in the quarantine listing to view the Manage by Rule Summary page. You can perform message actions (Release, Delete, Delay Exit) on all of the messages in the quarantine based on which outbreak rule caused the message to be quarantined. This is

ideal for clearing out large numbers of messages from the Outbreak quarantine. For more information, see information about the Manage by Rule Summary view in the Outbreak Filters chapter in the online help or user guide for the Email Security appliance

Reporting False Positives or Suspicious Messages to Cisco Systems

When viewing message details for a message in the Outbreak quarantine, you can send the message to Cisco to report false positives or suspicious messages.

Step 1 Navigate to a message in the Outbreak quarantine.

Step 2 Click on the check box of the message(s) and select Send a Copy  .

Step 3 Enter the recipient address and click Send.

Troubleshooting Centralized Policy Quarantines

Administrative User Cannot Choose Quarantines in Filters and DLP Message Actions

Problem

Administrative users cannot see or choose quarantines in content and message filters or DLP actions on the Email Security appliance.

Solution

See Configuring Centralized Quarantine Access for Custom User Roles , on page 205

Messages Released from a Centralized Outbreak Quarantine Are Not Rescanned

Problem

Messages released from the Outbreak Quarantine should be scanned again before delivery. However, some contaminated messages have been delivered from the quarantine.

Solution

This can occur under the situation described in About Rescanning of Quarantined Messages , on page 220

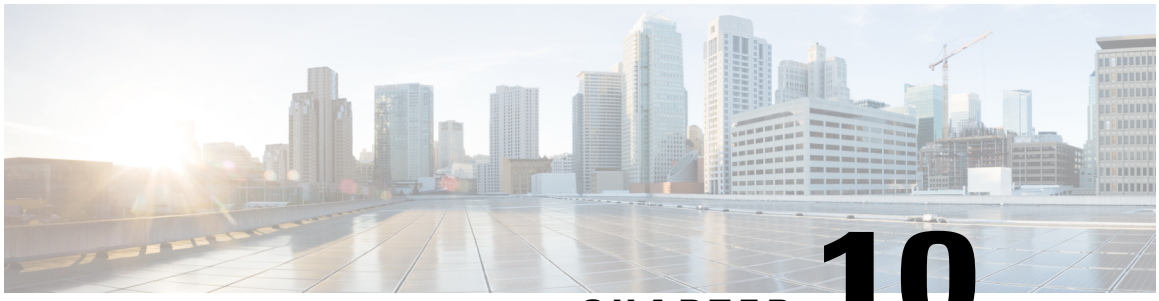

CHAPTER 10

Monitoring System Status

This chapter contains the following sections:

About Security Management Appliance Status

By default, the System Status page is the first page that appears when you access the Cisco Content Security Management appliance from your browser. (To change the landing page, see Setting Preferences , on page 346.)

To access the System Status page at any other time, select **Management Appliance > Centralized Services > System Status**.

Before you enable monitoring services and add a managed appliance, only the System Information section provides status information. If you have run the System Setup Wizard, enabled centralized services, and added managed appliances, the Centralized Services section and the Security Appliance Data Transfer Status section are populated with data.

Status information includes the following:

- **Centralized Services:**Status of each centralized service, including Processing Queue usage
- **System Uptime:** How long the appliance has been running
- **CPU Utilization:** Percentage of CPU capacity used by each monitoring service

- **System Version Information:** Model number, AsyncOS (operating system) version, build date, installation date, and serial number

Related Topics

- Monitoring the Processing Queue , on page 224
- Monitoring CPU Utilization , on page 225
- Monitoring Status of Data Transfer From Managed Appliances , on page 225

Monitoring Security Management Appliance Capacity

Monitoring the Processing Queue

You can periodically check the processing queue percentages used for email reporting and tracking to determine whether your appliance is running at optimal capacity.

The processing queue stores centralized reporting and tracking files as they await processing by the Security Management appliance. Normally, the Security Management appliance receives batches of reporting and tracking files for processing. The percentage of reporting or tracking files in the processing queue typically fluctuates as the files are transmitted from managed appliances and processed by the Security Management appliance.

Note Processing queue percentages gauge the number of files in the queue. They do not take file size into account. The percentages provide only a rough estimate of the Security Management appliance's processing load.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Centralized Services > System Status**.

Step 3 In the **Centralized Services** section at the top of the page, look at the Processing Queue percentages for:

a) Centralized Reporting (Email Security subsection)
b) Centralized Message Tracking

Step 4 If the processing queue usage percentages remain consistently high over several hours or days, then the system is running at or beyond capacity.
In that case, consider removing some of the managed appliances from the Security Management appliance, installing additional Security Management appliances, or both.

Monitoring CPU Utilization

To view the percentage of its CPU capacity that the Security Management appliance is using for each centralized service:

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Select **Management Appliance > Centralized Services > System Status**.

Step 3 Scroll to the **System Information** section and view the **CPU Utilization** subsection.
The CPU Utilization percentages indicate the portion of the Security Management appliance's CPU processing that is devoted to each of the main centralized services. Utilization percentages for some services may be combined. For example, email reporting are combined under "Reporting Service" while spam, policy, virus, and outbreak quarantines are combined under "Quarantine Services." Other operations of the Security Management appliance are grouped under the general heading "Security Management appliance."

Step 4 Refresh the browser display to view the most recent data.
The CPU utilization percentages change constantly.

Monitoring Status of Data Transfer From Managed Appliances

To perform centralized management functions, the Security Management appliance relies on the successful transfer of data from the managed appliances to the Security Management appliance. The Security Appliance Data Transfer Status section provides status information about each appliance that is managed by the Security Management appliance.

By default, the Security Appliance Data Transfer Status section displays up to ten appliances. If the Security Management appliance manages more than ten appliances, you can use the Items Displayed menu to select the number of appliances to display.

Note Summary information about data transfer status appears in the Services section at the top of the System Status page. The Security Appliance Data Transfer Status section provides appliance-specific data transfer status.

In the Security Appliance Data Transfer Status section of the System Status page, you can view connection status issues for specific appliances. For detailed information about the status of each service on an appliance, click the appliance name to view the Data Transfer Status page for the appliance.

The Data Transfer Status: *Appliance_Name* page shows when the last data transfer occurred for each monitoring service.

The data transfer status for Email Security appliances can be one of the following values:

- **Not enabled:**The monitoring service is not enabled on the Email Security appliance.
- **Never connected:** The monitoring service is enabled on the Email Security appliance, but no connection has been established between the Email Security appliance and the Security Management appliance.

- **Waiting for data:** The Email Security appliance has connected to the Security Management appliance, which is waiting to receive data.
- **Connected and transferred data:**A connection was established between the Email Security appliance and the Security Management appliance, and data were successfully transferred.
- **File transfer failure:** A connection was established between the Email Security appliance and the Security Management appliance, but the data transfer failed.

Data transfer issues can reflect temporary network problems or appliance configuration issues. The statuses of "Never connected" and "Waiting for data" are normal, transient statuses when you first add a managed appliance to the Security Management appliance. If the status does not eventually change to "Connected and transferred data," then the data transfer status might indicate a configuration issue.

If the "File transfer failure" status appears for an appliance, monitor the appliance to determine if the failure was caused by a network issue or by a problem with the appliance configuration. If no network issues prevent data transfer and the status does not change to "Connected and transferred data," then you might need to change the appliance configuration to enable data transfer.

Viewing the Configuration Status of Your Managed Appliances

On the Security Management appliance, choose **Management Appliance > Centralized Services > Security Appliances**.

The Centralized Service Status section shows which services are enabled and how many licenses you have used for each service. The Security Appliances section lists the appliances you have added. Check marks indicate the enabled services, and the Connection Established? column shows whether or not file transfer access is properly configured.

Related Topics

- Designating an Alternate Appliance to Process Released Messages, on page 205
- About Adding Managed Appliances, on page 19

Monitoring Reporting Data Availability Status

The Security Management appliance enables you to monitor the availability of reporting data for a specified time period. See the appropriate section for your appliance:

Monitoring Email Security Reporting Data Availability

To monitor reporting data from your Email Security appliances on the Security Management appliance, view the **Email > Reporting > Reporting Data Availability** page.

From the **Reporting Data Availability** page, you can view the percentage of reporting data that the Security Management appliance received from your Email Security appliances over a specified period of time. A bar chart indicates the completeness of the data received during the time range.

You can monitor reporting data availability for the preceding day, week, month, or year. If the Security Management appliance received less than 100% of the reporting data from the Email Security appliances, you

can tell immediately that your data may be incomplete. Use the data availability information to validate reporting data and to troubleshoot system problems.

Monitoring Email Tracking Data Status

To monitor the status of email tracking data, view the **Email > Message Tracking > Message Tracking Data Availability** page.

Identifying Active TCP/IP Services

To identify active TCP/IP services used by your Security Management appliance, use the tcpservices command in the command line interface.

Replacing a Managed Appliance During Hardware Failure

If you have to replace a managed appliance due to a hardware failure or other reasons, the data from the replaced appliance will not be lost, but the data will not be displayed correctly on the Security Management appliance.

Upon replacing a managed appliance, add the new appliance to the list of hosts on the SMA, and connect it to the new appliance. If the IP address remains the same, change the IP on the old host entry to a non existing value.

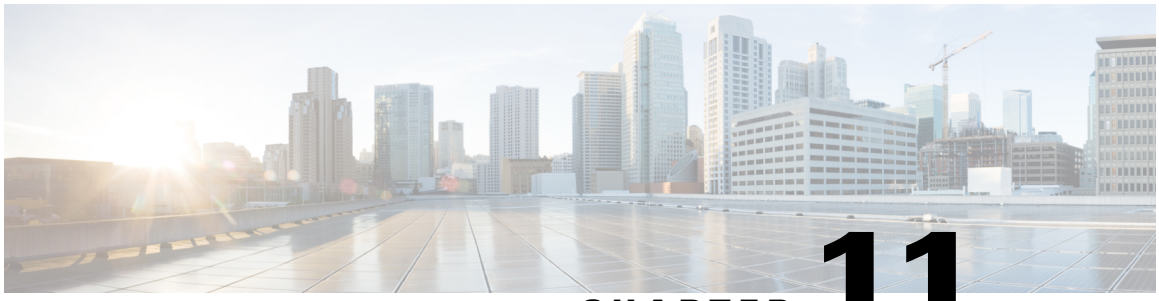

CHAPTER 11

Integrating With LDAP

This chapter contains the following sections:

Overview

If you maintain end-user passwords and email aliases in a corporate LDAP directory — for example, in Microsoft Active Directory, SunONE Directory Server, or OpenLDAP directories — you can use the LDAP directory to authenticate the following users:

- End users and administrative users who access the spam quarantine.

  When a user logs in to the web UI for the spam quarantine, the LDAP server validates the login name and password, and AsyncOS retrieves a list of the corresponding email aliases. Quarantined messages sent to any of the user's email aliases can appear in the spam quarantine, as long as the appliance does not rewrite them.

  See Configuring LDAP to Work with the Spam Quarantine, on page 230.

- Administrative users who sign in to the Cisco Content Security Management appliance when External Authentication is enabled and configured.

  See Configuring External Authentication of Administrative Users Using LDAP , on page 242.

Configuring LDAP to Work with the Spam Quarantine

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

Step 1 **Configure an LDAP server profile**.
The server profile contains information to enable AsyncOS to connect to the LDAP server, such as:

- Server name and port
- Base DN
- Authentication requirements for binding to the server

For more information about configuring a server profile, see Creating the LDAP Server Profile, on page 230.

When you create the LDAP server profile, you can configure AsyncOS to connect to multiple LDAP servers. For more information, see Configuring AsyncOS to Work With Multiple LDAP Servers, on page 239.

Step 2 **Configure the LDAP queries**.
You can either use the default spam quarantine queries generated for the LDAP server profile or create your own queries that are tailored to your particular LDAP implementation and schema. You then designate the active queries for spam notifications and end-user access to the quarantine.

For information about queries, see Configuring LDAP Queries, on page 232.

Step 3 **Enable LDAP end-user access and spam notifications for the spam quarantine**.
Enable LDAP end-user access to the spam quarantine to allow end-users to view and manage messages in their quarantine. You can also enable alias consolidation for spam notifications to prevent the user from receiving multiple notifications.

For more information, see Setting Up the Centralized Spam Quarantine , on page 174.

Creating the LDAP Server Profile

When you configure AsyncOS to use LDAP directories, you create an LDAP server profile to store the information about the LDAP server.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > LDAP**.

Step 3 Click **Add LDAP Server Profile.**

Step 4 Enter a name for the server profile in the **LDAP Server Profile Name** text field.

Step 5 Enter the host name for the LDAP server in the **Host Name(s)** text field.

You can enter multiple host names to configure the LDAP servers for failover or load-balancing. Separate multiple entries with commas. For more information, see Configuring AsyncOS to Work With Multiple LDAP Servers, on page 239.

Step 6 Select an authentication method. You can use anonymous authentication or specify a user name and password.

Note You need to configure LDAP authentication to view client user IDs instead of client IP addresses on reports. Without LDAP authentication the system can only refer to users by their IP address. Choose the **Use Password** radio button, and enter the User name and password. The user name will now be seen on the Internal Users Summary page.

Step 7 Select the LDAP server type: Active Directory, OpenLDAP, or Unknown or Other.

Step 8 Enter a port number.
The default port is `3268`. This is the default port for Active Directory that enables it to access the global catalog in a multi-server environment.

Step 9 Enter a base DN (distinguishing name) for the LDAP server.
If you authenticate with a user name and a password, the user name must include the full DN to the entry that contains the password. For example, a user with an email address of `joe@example.com` is a user of the marketing group. The entry for this user would look like the following entry:

```
uid=joe, ou=marketing, dc=example dc=com
```

Step 10 Under Advanced, select whether to use SSL when communicating with the LDAP server.

Step 11 Enter the cache time-to-live. This value represents the amount of time to retain caches.

Step 12 Enter the maximum number of retained cache entries.

Step 13 Enter a maximum number of simultaneous connections.
If you configure the LDAP server profile for load balancing, these connections are distributed among the listed LDAP servers. For example, if you configure 10 simultaneous connections and load balance the connections over three servers, AsyncOS creates 10 connections to each server, for a total of 30 connections. For more information, see Load Balancing, on page 241.

Note The maximum number of simultaneous connections includes LDAP connections used for LDAP queries. However, if you enable LDAP authentication for the spam quarantine, the appliance allows 20 additional connections for the end user quarantine for a total of 30 connections.

Step 14 Test the connection to the server by clicking the Test Server(s) button. If you specified multiple LDAP servers, they are all tested. The results of the test appear in the Connection Status field. For more information, see Testing LDAP Servers, on page 232.

Step 15 Create spam quarantine queries by selecting the check box and completing the fields.
You can configure the quarantine end-user authentication query to validate users when they log in to the end-user quarantine. You can configure the alias consolidation query so that end-users do not receive quarantine notices for each email alias. To use these queries, select the “Designate as the active query” check box. For more information, see Configuring LDAP Queries, on page 232.

Step 16 Test the spam quarantine queries by clicking the Test Query button.
Enter the test parameters and click Run Test. The results of the test appear in the Connection Status field. If you make any changes to the query definition or attributes, click **Update**.

Note If you have configured the LDAP server to allow binds with empty passwords, the query can pass the test with an empty password field.

Step 17 Submit and commit your changes.
Active Directory server configurations do not allow authentication through TLS with Windows 2000. This is a known issue with Active Directory. TLS authentication for Active Directory and Windows 2003 *does* work.

Note Although the number of server configurations is unlimited, you can configure only one end-user authentication query and one alias consolidation query per server.

Testing LDAP Servers

Use the Test Server(s) button on the Add/Edit LDAP Server Profile page (or the test subcommand of the ldapconfig command in the CLI) to test the connection to the LDAP server. AsyncOS displays a message stating whether the connection to the server port succeeded or failed. If you configured multiple LDAP servers, AsyncOS tests each server and displays individual results.

Configuring LDAP Queries

The following sections provide the default query strings and configuration details for each type of spam quarantine query:

- **Spam quarantine end-user authentication query.** For more information, see the Spam Quarantine End-User Authentication Queries, on page 233.
- **Spam quarantine alias consolidation query.** For more information, see Spam Quarantine Alias Consolidation Queries, on page 234.

To have the quarantine use an LDAP query for end-user access or spam notifications, select the "Designate as the active query" check box. You can designate one end-user authentication query to control quarantine access and one alias consolidation query for spam notifications. Any existing active queries are disabled. On the Security Management appliance, choose **Management Appliance > System Administration > LDAP** page, an asterisk (*) is displayed next to the active queries.

You can also specify a domain-based query or chain query as an active end-user access or spam notification query. For more information, see Domain-Based Queries, on page 236 and Chain Queries, on page 238.

Note Use the Test Query button on the LDAP page (or the **ldaptest** command) to verify that your queries return the expected results.

LDAP Query Syntax

Spaces are allowed in LDAP paths, and they do not need to be quoted. The CN and DC syntax is not case-sensitive.

Cn=First Last,oU=user,dc=domain,DC=COM

The variable names you enter for queries are case-sensitive and must match your LDAP implementation in order to work correctly. For example, entering **mailLocalAddress** at a prompt performs a different query than entering **maillocaladdress**.

Tokens

You can use the following tokens in your LDAP queries:

- {a} username@domainname
- {d} domain
- {dn} distinguished name
- {g} group name
- {u} user name
- {f} MAILFROM: address

Note The {f} token is valid in acceptance queries only.

For example, you might use the following query to accept mail for an Active Directory LDAP server:
(|(mail={a})(proxyAddresses=smtp:{a}))

Note We strongly recommend using the Test feature of the LDAP page (or the **test** subcommand of the **ldapconfig** command) to test all queries you construct and ensure that expected results are returned before you enable LDAP functionality on a listener. See the Testing LDAP Queries, on page 236 for more information.

Spam Quarantine End-User Authentication Queries

End-user authentication queries validate users when they log in to the spam quarantine. The token {u} specifies the user (it represents the user's login name). The token {a} specifies the user's email address. The LDAP query does not strip "SMTP:" from the email address; AsyncOS strips that portion of the address.

Based on the server type, AsyncOS uses one of the following default query strings for the end-user authentication query:

- **Active Directory:** (sAMAccountName={u})
- OpenLDAP: (uid={u})
- **Unknown or Other:** [Blank]

By default, the primary email attribute is **mail**. You can enter your own query and email attributes. To create the query in the CLI, use the isqauth subcommand of the **ldapconfig** command.

Note If you want users to log in with their full email addresses, use (mail=smtp:{a}) for the query string.

Sample Active Directory End-User Authentication Settings

This section shows sample settings for an Active Directory server and the end-user authentication query. This example uses password authentication for the Active Directory server, the default query string for end-user authentication for Active Directory servers, and the mail and proxyAddresses email attributes.

Table 44: Example LDAP Server and Spam Quarantine End-User Authentication Settings: Active Directory

Authentication Method	Use Password (Need to create a low-privilege user to bind for searching, or configure anonymous searching.)
Server Type	Active Directory
Port	3268
Base DN	[Blank]
Connection Protocol	[Blank]
Query String	(sAMAccountName={u})
Email Attribute(s)	mail,proxyAddresses

Sample OpenLDAP End-User Authentication Settings

This section shows sample settings for an OpenLDAP server and the end-user authentication query. This example uses anonymous authentication for the OpenLDAP server, the default query string for end-user authentication for OpenLDAP servers, and the mail and mailLocalAddress email attributes.

Table 45: Example LDAP Server and Spam Quarantine End-User Authentication Settings: OpenLDAP

Authentication Method	Anonymous
Server Type	OpenLDAP
Port	389
Base DN	[Blank] (Some older schemas will want to use a specific Base DN.)
Connection Protocol	[Blank]
Query String	(uid={u})
Email Attribute(s)	mail,mailLocalAddress

Spam Quarantine Alias Consolidation Queries

If you use spam notifications, the spam quarantine alias consolidation query consolidates the email aliases so that recipients do not receive quarantine notices for each alias. For example, a recipient might receive mail for the following email addresses: john@example.com, jsmith@example.com, and john.smith@example.com. When you use alias consolidation, the recipient receives a single spam notification at a chosen primary email address for messages sent to all of the user's aliases.

To consolidate messages to a primary email address, create a query to search for a recipient's alternate email aliases, and then enter the attribute for the recipient's primary email address in the Email Attribute field.

For Active Directory servers, the default query string (which may or may not be different for your deployment) is (|(proxyAddresses={a})(proxyAddresses=smtp:{a})) and the default email attribute is mail . For OpenLDAP servers, the default query string is (mail={a}) and the default email attribute is mail . You can define your own query and email attributes, including multiple attributes separated by commas. If you enter more than one email attribute, Cisco recommends entering a unique attribute that uses a single value, such as mail , as the first email attribute instead of an attribute with multiple values that can change, such as proxyAddresses .

To create the query in the CLI, use the isqalias subcommand of the ldapconfig command.

Sample Active Directory Alias Consolidation Settings

This section shows sample settings for an Active Directory server and the alias consolidation query. This example uses anonymous authentication for the Active Directory server, a query string for alias consolidation for Active Directory servers, and the mail email attribute.

Table 46: Example LDAP Server and Spam Quarantine Alias Consolidation Settings: Active Directory

Authentication Method	Anonymous
Server Type	Active Directory
Port	3268
Base DN	[Blank]
Connection Protocol	Use SSL
Query String	(|(mail={a})(mail=smtp:{a}))
Email Attribute	mail

Sample OpenLDAP Alias Consolidation Settings

This section shows sample settings for an OpenLDAP server and the alias consolidation query. This example uses anonymous authentication for the OpenLDAP server, a query string for alias consolidation for OpenLDAP servers, and the mail email attribute.

Table 47: Example LDAP Server and Spam Quarantine Alias Consolidation Settings: OpenLDAP

Authentication Method	Anonymous
Server Type	OpenLDAP
Port	389

Authentication Method	Anonymous
Base DN	[Blank] (Some older schemas will want to use a specific Base DN.)
Connection Protocol	Use SSL
Query String	(mail={a}))
Email Attribute	mail

Testing LDAP Queries

Use the Test Query button on the Add/Edit LDAP Server Profile page (or the ldaptest command in the CLI) to test your queries. AsyncOS displays details about each stage of the query connection test. For example, whether the first stage SMTP authorization succeeded or failed, and whether the BIND match returned a true or false result.

The ldaptest command is available as a batch command, for example:

ldaptest LDAP.isqalias foo@cisco.com

The variable names you enter for queries are *case-sensitive* and must match your LDAP implementation to work correctly. For example, entering mailLocalAddress for the email attribute performs a different query than entering maillocaladdress .

To test a query, you must enter the test parameters and click Run Test. The results appear in the Test Connection field. If an end-user authentication query succeeds, a result of "Success: Action: match positive" is displayed. For alias consolidation queries, a result of "Success: Action: alias consolidation" is displayed, along with the email address for the consolidated spam notifications. If a query fails, AsyncOS displays a reason for the failure, such as no matching LDAP records were found, or the matching record did not contain the email attribute. If you use multiple LDAP servers, the Cisco Content Security appliance tests the query on each LDAP server.

Domain-Based Queries

Domain-based queries are LDAP queries that are grouped by type and associated with a domain. You might want to use domain-based queries if different LDAP servers are associated with different domains, but you need to run queries for all your LDAP servers for end-user quarantine access. For example, a company called Bigfish owns the domains Bigfish.com, Redfish.com, and Bluefish.com, and it maintains a different LDAP server for employees associated with each domain. Bigfish can use a domain-based query to authenticate end-users against the LDAP directories of all three domains.

To use a domain-based query to control end-user access or notifications for the spam quarantine, complete the following steps:

Step 1 Create an LDAP server profile for each domain you want to use in the domain-based query. In each server profile, configure the queries you want to use in the domain-based query. For more information, see Creating the LDAP Server Profile, on page 230.

Step 2 Create the domain-based query. When you create the domain-based query, you select queries from each server profile, and designate the domain-based query as an active query for the spam quarantine. For more information about creating the query, see Creating a Domain-Based Query, on page 237.

Step 3 Enable end-user access or spam notifications for the spam quarantine. For more information, see Setting Up End-User Access to the Spam Quarantine via Web Browser , on page 189.

Creating a Domain-Based Query

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > LDAP**.

Step 3 On the LDAP page, click **Advanced**.

Step 4 Enter a name for the domain-based query.

Step 5 Select the query type.

Note When you create a domain-based query, you specify a single query type. After you select a query type, the query field drop-down lists contain the appropriate queries from the LDAP server profiles.

Step 6 In the Domain Assignments field, enter a domain.

Step 7 Select a query to associate with the domain.

Step 8 Add a row and select a query for each domain in the domain-based query.

Step 9 Enter a default query to run if all other queries fail. If you do not want to enter a default query, select **None**.

Figure 5: Example Domain-based Query

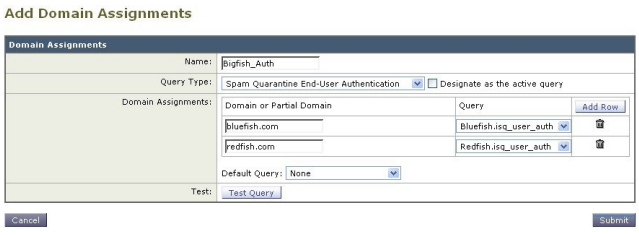

Step 10 Test the query by clicking the Test Query button and entering a user login and password or an email address to test in the Test Parameters fields. The results appear in the Connection Status field.

Step 11 Check the **Designate as the active query** checkbox if you want the spam quarantine to use the domain-based query.

Note The domain-based query becomes the active LDAP query for the specified query type. For example, if the domain-based query is used for end-user authentication, it becomes the active end-user authentication query for the spam quarantine.

Step 12 Click **Submit** and then click **Commit** to commit your changes.

Note To do the same configuration on the command line interface, type the `advanced` subcommand of the `ldapconfig` command at the command line prompt.

Chain Queries

A chain query is a series of LDAP queries that AsyncOS runs in succession. AsyncOS runs each query in the series each query in the "chain" until the LDAP server returns a positive response or the final query returns a negative response or fails. Chain queries can be useful if entries in LDAP directories use different attributes to store similar (or the same) values. For example, departments in an organization might use different types of LDAP directories. The IT department might use OpenLDAP while the Sales department uses Active Directory. To ensure that queries run against both types of LDAP directories, you can use chain queries.

To use a chain query to control end-user access or notifications for the spam quarantine, complete the following steps:

Step 1 Create an LDAP server profile for each query you want to use in the chain queries. For each of the server profiles, configure the queries you want to use for a chain query. For more information, see Creating the LDAP Server Profile, on page 230.

Step 2 Create the chain query and designate it as an active query for the spam quarantine. For more information, see Creating a Chain Query, on page 238.

Step 3 Enable LDAP end-user access or spam notifications for the spam quarantine. For more information about the spam quarantine, see Setting Up the Centralized Spam Quarantine , on page 174.

Creating a Chain Query

Tip You can also use the advanced subcommand of the ldapconfig command in the CLI.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > LDAP > LDAP Server**.

Step 3 From the LDAP Server Profiles page, click **Advanced**.

Step 4 Click **Add Chained Query**.

Step 5 Enter a name for the chain query.

Step 6 Select the query type.

When you create a chain query, all of its component queries have the same query type. After you select a query type, the query field drop-down lists display the appropriate queries from the LDAP.

Step 7 Select the first query in the chain.
The Cisco Content Security appliance runs the queries in the order you configure them. If you add multiple queries to the chain query, you might want to order them so that general queries follow granular queries.

Figure 6: Example Chain Query

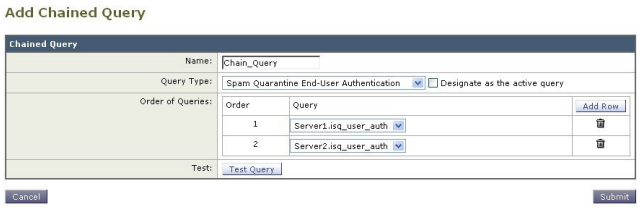

Step 8 Test the query by clicking the Test Query button and entering a user login and password or an email address in the Test Parameters fields. The results appear in the Connection Status field.

Step 9 Check the **Designate as the active query** check box if you want the spam quarantine to use the domain query.

Note The chain query becomes the active LDAP query for the specified query type. For example, if the chain query is used for end-user authentication, it becomes the active end-user authentication query for the spam quarantine.

Step 10 Submit and commit your changes.

Note To do the same configuration on the command line interface, type the `advanced` subcommand of the `ldapconfig` command at the command line prompt.

Configuring AsyncOS to Work With Multiple LDAP Servers

When you configure an LDAP server profile, you can configure the Cisco Content Security appliance to connect to a list of multiple LDAP servers. If you use multiple LDAP servers, they need to contain the same information, have the same structure, and use the same authentication information. Third-party products exist that can consolidate the records.

You configure the Cisco Content Security appliance to connect to redundant LDAP servers to use the following features:

- **Failover**. If the Cisco Content Security appliance cannot connect to an LDAP server, it connects to the next server in the list.
- **Load Balancing**. The Cisco Content Security appliance distributes connections across the list of LDAP servers when it performs LDAP queries.

You can configure redundant LDAP servers on the Management Appliance > System Administration > LDAP page or by using the CLI ldapconfig command.

Testing Servers and Queries

Use the Test Server(s) button on the Add (or Edit) LDAP Server Profile page (or the test subcommand in the CLI) to test the connection to an LDAP server. If you use multiple LDAP servers, AsyncOS tests each server

and displays individual results for each server. AsyncOS will also test the query on each LDAP server and display the individual results.

Failover

To ensure an LDAP server is available to that resolve queries, you can configure the LDAP profile for failover. If the connection to the LDAP server fails, or the query returns an error for which it is appropriate to do so, the appliance attempts to query the next LDAP server specified in the list.

The Cisco Content Security appliance attempts to connect to the first server in the list of LDAP servers for a specified period of time. If the appliance cannot connect to the first LDAP server in the list, or the query returns an error, the appliance attempts to connect to the next LDAP server in the list. By default, the appliance always attempts to connect to the first server in the list, and it attempts to connect to each subsequent server in the order they are listed. To ensure that the Cisco Content Security appliance connects to the primary LDAP server by default, enter it as the first server in the list of LDAP servers.

Note Only attempts to query a specified LDAP server fail over. Attempts to query referral or continuation servers associated with the specified LDAP server do not fail over.

If the Cisco Content Security appliance connects to a second or subsequent LDAP server, it remains connected to that server for a specified period of time. At the end of this period, the appliance attempts to reconnect to the first server in the list.

Configuring the Cisco Content Security Appliance for LDAP Failover

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > LDAP**.

Step 3 Select the LDAP server profile you want to edit.
In the following example, the LDAP server name is example.com.

Figure 7: Example LDAP Failover Configuration

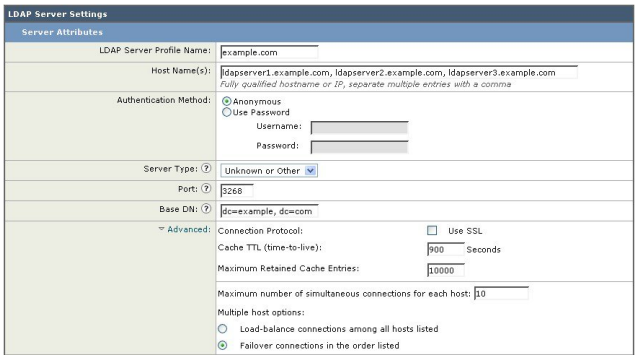

Step 4 In the Hostname text field, type the LDAP Servers; for example **ldapserver.example.com**.

Step 5 In the Maximum number of simultaneous connections for each host text field, type the maximum number of connections. In this example the maximum number of connections is **10**.

Step 6 Click on the radio button next to **Failover connections in the order list**.

Step 7 Configure other LDAP options as necessary.

Step 8 Submit and commit the changes.

Load Balancing

To distribute LDAP connections among a group of LDAP servers, you can configure your LDAP profile for load balancing.

When you use load balancing, the Cisco Content Security appliance distributes connections among the LDAP servers listed. If a connection fails or times out, the appliance determines which LDAP servers are available and reconnects to available servers. The appliance determines the number of simultaneous connections to establish based on the maximum number of connections you configure.

If one of the listed LDAP servers does not respond, the appliance distributes the connection load among the remaining LDAP servers.

Configuring the Cisco Content Security Appliance for Load Balancing

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > LDAP**.

Step 3 Select the LDAP server profile you want to edit
In the following example, the LDAP server name is example.com.

Figure 8: Example Loadbalancing Configuration

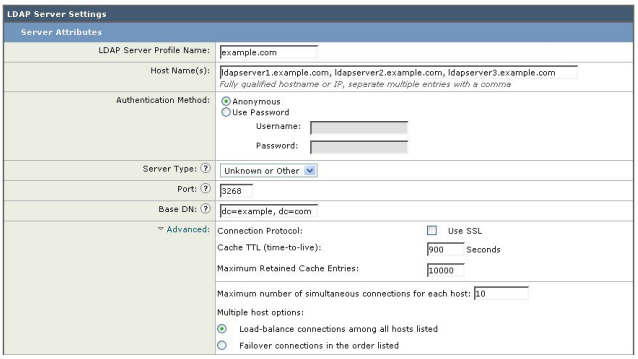

Step 4 In the Hostname text field, type the LDAP Servers; for example **ldapserver.example.com**.

Step 5 In the Maximum number of simultaneous connections for each host text field, type the maximum number of connections. In this example the maximum number of connections is **10**.

Step 6 Click on the radio button next to **Load balance connections among all hosts**.

Step 7 Configure other LDAP options as necessary.

Step 8 Submit and commit the changes.

Configuring External Authentication of Administrative Users Using LDAP

You can configure the Cisco Content Security appliance to use an LDAP directory on your network to authenticate administrative users by allowing them to log in to the appliance with their LDAP user names and passwords.

Step 1 **Configure the LDAP Server Profile.** See Creating the LDAP Server Profile, on page 230.

Step 2 **Create a query to find user accounts.** In an LDAP server profile, in the External Authentication Queries section, create a query to search for user accounts in the LDAP directory. See User Accounts Query for Authenticating Administrative Users , on page 242.

Step 3 **Create group membership queries.** Create a query to determine if a user is a member of a directory group, and create a separate query to find all members of a group. For more information, see Group Membership Queries for Authenticating Administrative Users , on page 243 and the documentation or online help for your Email Security appliance.

Note Use the **Test Queries** button in the External Authentication Queries section of the page (or the ldaptest command) to verify that your queries return the expected results. For related information, see Testing LDAP Queries, on page 236.

Step 4 **Set up external authentication to use the LDAP server.** Enable the appliance to use the LDAP server for user authentication and assign user roles to the groups in the LDAP directory. For more information, see Enabling External Authentication of Administrative Users , on page 245 and the “Adding Users” in the documentation or online help for your Email Security appliance.

User Accounts Query for Authenticating Administrative Users

To authenticate external users, AsyncOS uses a query to search for the user record in the LDAP directory and the attribute that contains the user’s full name. Depending on the server type you select, AsyncOS enters a default query and a default attribute. You can choose to have your appliance deny users with expired accounts if you have attributes defined in RFC 2307 in your LDAP user records (**shadowLastChange**, **shadowMax**, and **shadowExpire**). The base DN is required for the domain level where user records reside.

The following table shows the default query string and full user name attribute that AsyncOS uses when it searches for a user account on an Active Directory server.

Table 48: Default Query String for Active Directory Server

Server Type	Active Directory
Base DN	[blank] (You need to use a specific base DN to find the user records.)
Query String	(&(objectClass=user)(sAMAccountName={u}))
Attribute containing the user's full name	displayName

The following table shows the default query string and full user name attribute that AsyncOS uses when it searches for a user account on an OpenLDAP server.

Table 49: Default Query String for Open LDAP Server

Server Type	OpenLDAP
Base DN	[blank] (You need to use a specific base DN to find the user records.)
Query String	(&(objectClass=posixAccount)(uid={u}))
Attribute containing the user's full name	gecos

Group Membership Queries for Authenticating Administrative Users

You can associate LDAP groups with user roles for accessing the appliance.

AsyncOS also uses a query to determine if a user is a member of a directory group and a separate query to find all members of a group. Membership in a directory group membership determines the user's permissions within the system. When you enable external authentication on the Management Appliance > System Administration > Users page in the GUI (or userconfig in the CLI), you assign user roles to the groups in your LDAP directory. User roles determine the permissions that users have in the system, and for externally authenticated users, the roles are assigned to directory groups instead of individual users. For example, you can assign users in the IT directory group the Administrator role and users in the Support directory group to the Help Desk User role.

If a user belongs to multiple LDAP groups with different user roles, AsyncOS grants the user the permissions for the most restrictive role. For example, if a user belongs to a group with Operator permissions and a group with Help Desk User permissions, AsyncOS grants the user the permissions for the Help Desk User role.

When you configure the LDAP profile to query for group membership, enter the base DN for the directory level where group records can be found, the attribute that holds the group member's user name, and the attribute that contains the group name. Based on the server type that you select for your LDAP server profile, AsyncOS enters default values for the user name and group name attributes, as well default query strings.

Note For Active Directory servers, the default query string to determine if a user is a member of a group is (&(objectClass=group)(member={u})). However, if your LDAP schema uses distinguished names in the "memberof" list instead of user names, you can use {dn} instead of {u}.

The following table shows the default query strings and attributes that AsyncOS uses when it searches for group membership information on an Active Directory server.

Table 50: Default Query String and Attributes for Active Directory Server

Query String	Active Directory
Base DN	[blank] (You need to use a specific base DN to find the group records.)
Query string to determine if a user is a member of a group	(&(objectClass=group)(member={u})) **Note** If your LDAP schema uses distinguished names in the member of list instead of user names, you can replace {u} with {dn}
Query string to determine all members of a group	(&(objectClass=group)(cn={g}))
Attribute that holds each member's user name (or a DN for the user's record)	member
Attribute that contains the group name	cn

The following table shows the default query strings and attributes that AsyncOS uses when it searches for group membership information on an OpenLDAP server.

Table 51: Default Query String and Attributes for Open LDAP Server

Query String	OpenLDAP
Base DN	[blank] (You need to use a specific base DN to find the group records.)
Query string to determine if a user is a member of a group	(&(objectClass=posixGroup)(memberUid={u}))
Query string to determine all members of a group	(&(objectClass=posixGroup)(cn={g}))
Attribute that holds each member's user name (or a DN for the user's record)	memberUid
Attribute that contains the group name	cn

Enabling External Authentication of Administrative Users

After you configure the LDAP server profile and queries, you can enable external authentication using LDAP:

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > Users** page.

Step 3 Click **Enable**.

Step 4 Select the **Enable External Authentication** check box.

Step 5 Select **LDAP** for the authentication type.

Step 6 Select the LDAP external authentication query that authenticates users.

Step 7 Enter the number of seconds that the appliance waits for a response from the server before timing out.

Step 8 Enter the name of a group from the LDAP directory that you want the appliance to authenticate, and select the role for the users in the group.

Step 9 Optionally, click **Add Row** to add another directory group. Repeat steps 7 and 8 for each directory group that the appliance authenticates.

Step 10 Submit and commit your changes.

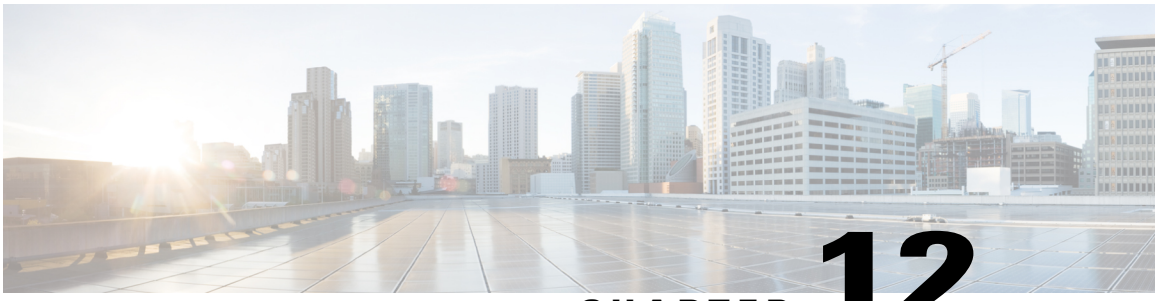

CHAPTER 12

Configuring SMTP Routing

This chapter contains the following sections:

SMTP Routes Overview

This chapter explains the features that affect routing and delivery of email traveling through the Cisco Content Security Management appliance, and use of the SMTP Routes page and **smtproutes** command.

SMTP Routes allow you to redirect all email for a particular domain to a different mail exchange (MX) host. For example, you could make a mapping from example.com to groupware.example.com . This mapping causes any email with @example.com in the Envelope Recipient address to go instead to groupware.example.com . The system performs an “MX” lookup on groupware.example.com , and then performs an “A” lookup on the host, just like a normal email delivery. This alternate MX host does not need to be listed in DNS MX records and it does not even need to be a member of the domain whose email is being redirected. The operating system allows up to ten thousand (10,000) SMTP Route mappings to be configured for your Cisco Content Security appliance. (See SMTP Routes Limits, on page 249.)

This feature also allows host “globbing.” If you specify a partial domain, such as example.com , then any domain ending in example.com matches the entry. For instance, fred@foo.example.com and wilma@bar.example.com both match the mapping.

If a host is not found in the SMTP Routes table, an MX lookup is performed using DNS. The result is not re-checked against the SMTP Routes table. If the DNS MX entry for foo.domain is bar.domain , any email sent to foo.domain is delivered to the host bar.domain . If you create a mapping for bar.domain to some other host, email addressed to foo.domain is not affected.

In other words, recursive entries are not followed. If there is an entry for a.domain to redirect to b.domain , and a subsequent entry to redirect email for b.domain to a.domain , a mail loop will *not* be created. In this case, email addressed to a.domain will be delivered to the MX host specified by b.domain , and conversely email addressed to b.domain will be delivered to the MX host specified by a.domain .

The SMTP Routes table is read from the top down for every email delivery. The most specific entry that matches a mapping wins. For example, if there are mappings for both host1.example.com and example.com

in the SMTP Routes table, the entry for host1.example.com will be used because it is the more specific entry — even if it appears after the less specific example.com entry. Otherwise, the system performs a regular MX lookup on the domain of the Envelope Recipient.

SMTP Routes, Mail Delivery, and Message Splintering

Incoming: if one message has 10 recipients and they are all on the same Exchange server, AsyncOS will open one TCP connection and present exactly one message to the mail store, not 10 separate messages.

Outgoing: works similarly, but if one message is going to 10 recipients in 10 different domains, AsyncOS will open 10 connections to 10 MTAs and deliver them one email each.

Splintering: if one incoming message has 10 recipients and they are each in separate Incoming Policy groups (10 groups), the message will splinter even if all 10 recipients are on the same Exchange server. Thus, 10 separate emails will be delivered over a single TCP connection.

SMTP Routes and Outbound SMTP Authentication

If an Outbound SMTP Authentication profile has been created, you can apply it to an SMTP Route. This allows authentication for outgoing mail in cases where the Cisco Content Security appliance sits behind a mail relay server that is at the edge of the network.

Routing Email for Local Domains

The Security Management appliance routes the following mail:

- ISQ released messages which ignore SMTP routing
- Alerts
- Configuration files that can be mailed to the specified destination
- Support request message that can be sent to the defined recipient as well

The last two types of messages use SMTP routes to be delivered to the destination.

The Email Security appliance routes mail to local domains to hosts specified using the **Management Appliance > Network > SMTP Routes** page (or the **smtproutes** command). This feature is similar to the sendmail **mailertable** feature. (The SMTP Routes page and **smtproutes** command are an expansion of the AsyncOS 2.0 Domain Redirect feature.)

Note If you have completed the System Setup Wizard in the GUI and committed the changes, you defined the first SMTP route entries on the appliance for each RAT entry you entered at that time.

Default SMTP Route

You can also define a default SMTP route with the special keyword ALL . If a domain does not match a previous mapping in the SMTP Routes list, it defaults to being redirected to the MX host specified by the ALL entry.

When you print the SMTP Routes entries, the default SMTP route is listed as ALL: . You cannot delete the default SMTP route; you may only clear any values entered for it.

Configure the default SMTP route using the **Management Appliance > Network > SMTP Routes** page or the **smtproutes** command.

Managing SMTP Routes

Defining an SMTP Route

The Email Security appliance routes mail to local domains to hosts specified using the **Management Appliance > Network > SMTP Routes** page (or the **smtproutes** command). This feature is similar to the sendmail mailer table feature. (The SMTP Routes page and **smtproutes** command are an expansion of the AsyncOS 2.0 Domain Redirect feature.):

Use the Management Appliance > Network > SMTP Routes page (or the **smtproutes** command) to construct routes. When you create a new route, you first specify the domain or partial domain for which you want to create a permanent route. You then specify destination hosts. Destination hosts can be entered as fully-qualified hostnames or as IP addresses. You can also specify a a special destination host of /dev/null to drop the messages that match the entry. (So, in effect, specifying /dev/null for the default route is will ensure that no mail received by the appliance is ever delivered.)

Multiple destination host entries can contain both fully-qualified hostnames and IP addresses. Separate multiple entries with commas.

If one or more of the hosts are not responding, messages will be delivered to one of the reachable hosts. If all the configured hosts are not responding, mail will be queued for that host (does not fail over to using MX records).

SMTP Routes Limits

You can define up to 10,000 routes. The final default route of ALL is counted as a route against this limit. Therefore, you can define up to 9,999 custom routes and one route that uses the special keyword ALL .

Adding SMTP Routes

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Network > SMTP Routes**.

Step 3 Click **Add Route**.

Step 4 Enter a receiving domain and destination host. You can add multiple destination hosts by clicking **Add Row** and entering the next destination host in the new row.

Step 5 You can specify a port number by adding ":<*port number*> " to the destination host: example.com:25

Step 6 Submit and commit your changes.

Exporting SMTP Routes

Similar to the Host Access Table (HAT) and the Recipient Access Table (RAT), you can also modify SMTP routes mappings by exporting and importing a file.

Step 1 Click **Export SMTP Routes** on the SMTP Routes page.

Step 2 Enter a name for the file and click **Submit**.

Importing SMTP Routes

Similar to the Host Access Table (HAT) and the Recipient Access Table (RAT), you can also modify SMTP routes mappings by exporting and importing a file.

Step 1 Click **Import SMTP Routes** on the SMTP Routes page.

Step 2 Select the file that contains the exported SMTP Routes.

Step 3 Click **Submit**. You are warned that importing will replace all existing SMTP Routes. All of the SMTP Routes in the text file are imported.

Step 4 Click **Import**.
You can place “comments” in the file. Lines that begin with a ‘#’ character are considered comments and are ignored by AsyncOS. For example:

this is a comment, but the next line is not

ALL:

At this point, our Email Gateway configuration looks like this:

Figure 9: Email Gateway Configuration

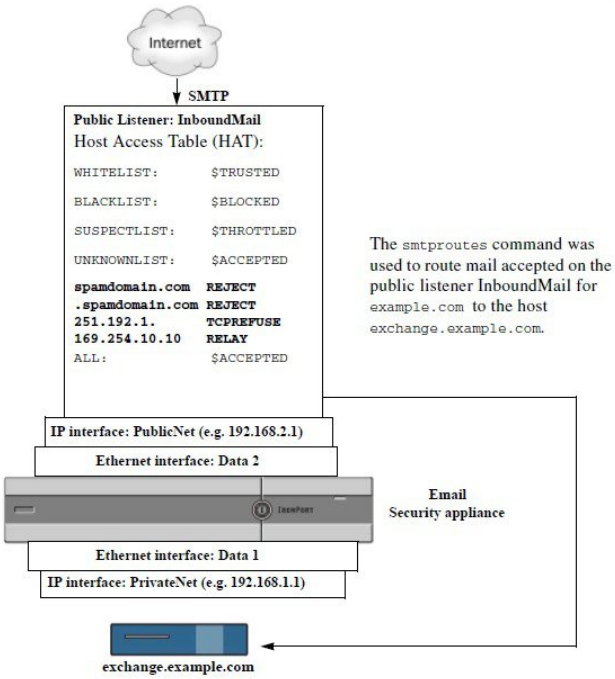

SMTP Routes and DNS

Use the special keyword USEDNS to tell the appliance to do MX lookups to determine next hops for specific domains. This is useful when you need to route mail for subdomains to a specific host. For example, if mail to example.com is to be sent to the company's Exchange server, you might have something similar to the following SMTP route:

```
example.com exchange.example.com
```

However, for mail to various subdomains (foo.example.com), add an SMTP route that looks like this:

```
.example.com USEDNS
```

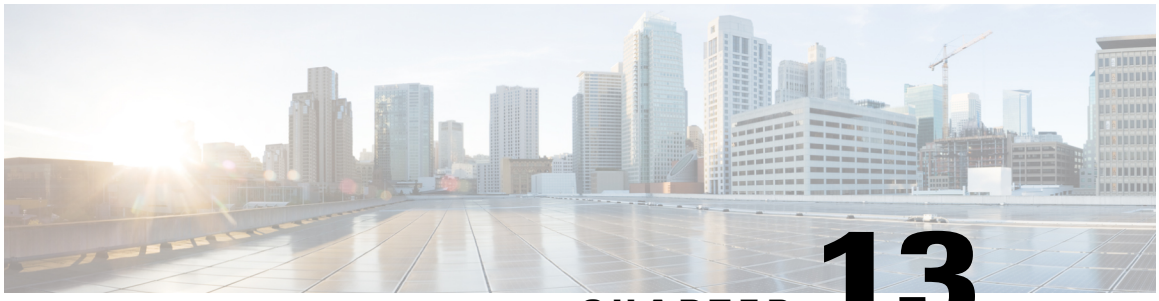

CHAPTER 13

Distributing Administrative Tasks

This chapter contains the following sections:

About Distributing Administrative Tasks

You can distribute administrative tasks on the Cisco Content Security Management appliance to other people based on the user roles that you assign to their user accounts.

To set up to distribute administrative tasks, you will determine whether the predefined user roles meet your needs, create any needed custom user roles, and set up the appliance to authenticate administrative users locally on the security appliance, and/or externally using your own centralized LDAP or RADIUS system.

Additionally, you can specify additional controls on access to the appliance and to certain information on the appliance.

Assigning User Roles

Additional configuration is required for quarantine access. See Access to Quarantines , on page 260.

Predefined User Roles

Except as noted, you can assign each user a predefined user role with the privileges described in the following table, or a custom user role.

Table 52: Descriptions of User Roles

User Role Name	Description	Scheduled Reports Capability
admin	The **admin** user is the default user account for the system and has all administrative privileges. The admin user account is listed here for convenience, but it cannot be assigned via a user role, and it cannot be edited or deleted, aside from changing the password. Only the **admin** user can issue the **resetconfig** and **revert**commands.	Yes
Administrator	User accounts with the Administrator role have full access to all configuration settings of the system.	Yes
Operator	User accounts with the Operator role are restricted from: • Creating or editing user accounts • Upgrading the appliance • Issuing the resetconfig command • Running the System Setup Wizard • Modifying LDAP server profile settings other than username and password, if LDAP is enabled for external authentication. • Configuring, editing, deleting, or centralizing quarantines. Otherwise, they have the same privileges as the Administrator role.	Yes
Technician	User accounts with the Technician role can initiate system administration activities such as upgrades and reboots, save a configuration file from the appliance, manage feature keys, and so forth.	Access to System Capacity reports under the Email tabs

User Role Name	Description	Scheduled Reports Capability
Read-Only Operator	User accounts with the Read-Only Operator role have access to view configuration information. Users with the Read-Only Operator role can make and submit most changes to see how to configure a feature, but they cannot commit them or make any change that does not require a commit. Users with this role can manage messages in quarantines, if access is enabled. Users with this role cannot access the following: • File system, FTP, or SCP. • Settings for creating, editing, deleting or centralizing quarantines.	No
Guest	Users accounts with the Guest role can view status information including reports and Tracking, and manage messages in quarantines, if access is enabled. Users with the Guest role cannot access Message Tracking.	No
Email Administrator	User accounts with the Email Administrator role have access to all configuration settings within the Email menu only, including quarantines.	No
Help Desk User	User accounts with the Help Desk User role are restricted to: • Message Tracking • Managing messages in quarantines Users with this role cannot access the rest of the system, including the CLI. After you assign a user this role, you must also configure quarantines to allow access by this user.	No

User Role Name	Description	Scheduled Reports Capability
Custom Roles	User accounts that are assigned a custom user role can view and configure only policies, features, or specific policy or feature instances that have been specifically delegated to the role. You can create a new Custom Email User Role from the Add Local User page. However, you must assign privileges to this Custom User Role before the role can be used. To assign privileges, go to Management Appliance > System Administration > User Roles and click the user name. **Note** Users assigned to a Custom Email User Role cannot access the CLI. For more information, see Custom User Roles , on page 256.	No

Custom User Roles

The Security Management appliance allows users with Administration privileges to delegate administration capabilities to custom roles. Custom roles provide more flexible control over your users' access than the predefined user roles do.

Users to whom you assign custom user roles can manage policies or access reports for a subset of appliances, features, or end users. For example, you might allow a delegated administrator for web services to manage policies for an organization's branch office in a different country, where the acceptable use policies might be different from those at the organization's headquarters. You delegate administration by creating custom user roles and assigning access permissions to those roles. You determine which policies, features, reports, custom URL categories, etc. that the delegated administrators can view and edit.

For more information, see:

- About Custom Email User Roles , on page 256
- Deleting Custom User Roles , on page 259

About Custom Email User Roles

You can assign custom roles to allow delegated administrators to access the following on the Security Management appliance:

- All reports (optionally restricted by Reporting Group)
- Mail Policy reports (optionally restricted by Reporting Group)
- DLP reports (optionally restricted by Reporting Group)
- Message Tracking
- Quarantines

Detailed information about each of these items follows this section. In addition, all users granted any of these privileges can see the System Status, available under the Management Appliance tab > Centralized Services menu. Users assigned to custom email user roles cannot access the CLI.

Note Custom user roles on the Email Security appliance offer more granular access than do user roles on the Security Management appliance. For example, you can delegate access to mail and DLP policies and content filters. For details, see the "Managing Custom User Roles for Delegated Administration" section in the "Common Administration" chapter of the documentation or online help for your Email Security appliance.

Access to Email Reporting

You can grant custom user roles access to Email reports as described in the following sections.

For complete information about the Email Security Monitor pages on the Security Management appliance, see the chapter on Using Centralized Email Security Reporting on the Legacy Web Interface, on page 35.

All Reports

If you grant a custom role access to All Reports, users assigned to this role can see the following Email Security Monitor pages either for all Email Security appliances, or for the Reporting Group that you select:

- Mail Flow Summary
- Mail FLow Details
- Outgoing Destinations
- User Mail Summary
- DLP Incidents
- Content Filters
- Virus Filtering
- TLS Encryption
- Scheduled Reports
- Archived Reports

Mail Policy Reports

If you grant a custom role access to Mail Policy Reports, users assigned to this role can see the following Email Security Monitor pages either for all Email Security appliances, or for the Reporting Group that you select:

- Mail Flow Summary
- Mail FLow Details
- Outgoing Destinations
- User Mail Summary
- Content Filters

- Virus Filtering
- Archived Reports

DLP Reports

If you grant a custom role access to DLP Reports, users assigned to this role can see the following Email Security Monitor pages either for all Email Security appliances, or for the Reporting Group that you select:

- DLP Incidents
- Archived Reports

Access to Message Tracking Data

If you grant a custom role access to Message Tracking, users to whom you assign this role can find the status of all messages tracked by the Security Management appliance.

To control access to sensitive information in messages that violate DLP policies, see Controlling Access to Sensitive Information in Message Tracking , on page 276.

For more information about message tracking, including instructions for setting up your appliances to enable access to message tracking on the Security Management appliance, see Tracking Messages, on page 159.

Access to Quarantines for Custom User Role

If you grant a custom role access to quarantines, users to whom you assign this role can search for, view, release, or delete messages in all quarantines on this Security Management appliance.

Before users can access quarantines, you must enable that access. See Access to Quarantines , on page 260.

Creating Custom Email User Roles

You can create custom email user roles for access to Email Reporting, Message Tracking, and quarantines.

For descriptions of the access that each of these options permits, see About Custom Email User Roles , on page 256 and its subsections.

Note To grant more granular access or access to other features, reports, or policies, create custom user roles directly on each Email Security appliance.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > User Roles**.

Step 3 Click **Add Email User Role**.

Tip Alternatively, you can create a new role by duplicating an existing Email User Role: Click the Duplicate icon in the applicable table row, then modify the resulting copy.

Step 4 Enter a unique name for the user role (for example, “dlp-auditor”) and a description.

- Email and Web custom user role names must not be duplicated.

- The name must contain only lowercase letters, numbers, and dashes. It cannot start with a dash or a number.
- If you grant users with this role access to centralized policy quarantines, and you also want users with this role to be able to specify those centralized quarantines in message and content filters and DLP Message Actions on an Email Security appliance, the name of the custom role must be the same on both appliances.

Step 5 Choose the access privileges to enable for this role.

Step 6 Click **Submit** to return to the User Roles page, which lists the new user role.

Step 7 If you limited access by Reporting Group, click the **no groups selected** link in the Email Reporting column for the user role, then choose at least one Reporting Group.

Step 8 Commit your changes.

Step 9 If you granted this role access to quarantines, enable access for this role:
See:

- Configuring Administrative User Access to the Spam Quarantine , on page 179
- Configuring Policy, Virus, and Outbreak Quarantines , on page 208

Using Custom Email User Roles

When a user who is assigned a custom email user role logs into the appliance, that user sees only the links to the security features to which that user has access. The user can return to this main page at any time by selecting Account Privileges in the Options menu. These users can also access the features to which they have access by using the menus at the top of the web page. In the following example, the user has access to all features that are available on the Security Management appliance via custom email user roles.

Figure 10: Account Privileges Page for a Delegated Administrator assigned Custom Email User Roles

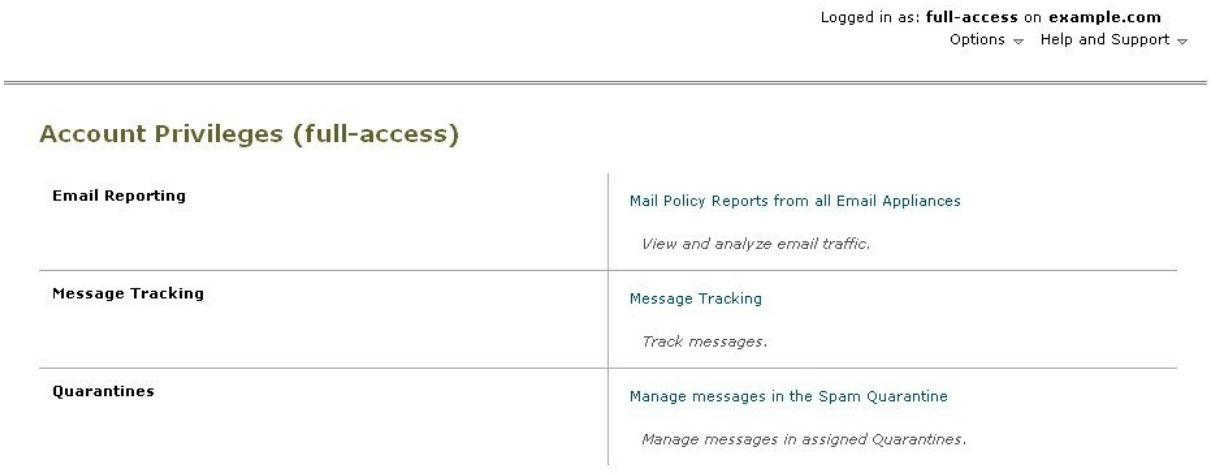

Deleting Custom User Roles

If you delete a custom user role that is assigned to one or more users, you do not receive an error.

User Roles with Access to the CLI

Some roles can access both the GUI and the CLI: Administrator, Operator, Guest, Technician, and Read-Only Operator. Other roles can access the GUI only: Help Desk User, Email Administrator,and custom user.

Using LDAP

If you use an LDAP directory to authenticate users, you assign directory groups to user roles instead of to individual users. When you assign a directory group to a user role, each user in that group receives the permissions defined for the user role. For more information, see External User Authentication , on page 268.

Access to Quarantines

Before users can access quarantines, you must enable that access. See the following information:

- Configuring Administrative User Access to the Spam Quarantine , on page 179
- About Distributing Message Processing Tasks to Other Users , on page 214 (for policy quarantines), and Configuring Policy, Virus, and Outbreak Quarantines , on page 208
- Configuring Centralized Quarantine Access for Custom User Roles , on page 205.

Users Page

For Information About This Section	See
Users Reset Passwords button	About Distributing Administrative Tasks , on page 253 Managing Locally-Defined Administrative Users , on page 261 Requiring Users to Change Passwords on Demand , on page 267
Local User Account & Password Settings	Setting Password and Login Requirements , on page 263
External Authentication	External User Authentication , on page 268
DLP Tracking Privileges	Controlling Access to Sensitive Information in Message Tracking , on page 276

About Authenticating Administrative Users

You can control access to the appliance by defining authorized users locally on the appliance, and/or by using external authentication or two-factor authentication.

Changing the Admin User's Password

Any administrator-level user can change the password for the "admin" user, via the GUI or the CLI.

To change the password via the GUI, do the following:

- [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.
- Choose **Management Appliance > System Administration > Users** page and select the admin user.

To change the password for the admin user in the CLI, use the `password` command. The `password` command requires you to enter the old password for security.

If you forget the password for the "admin" user account, contact your customer support provider to reset the password.

Note Changes to the password take effect immediately and do not require you to commit the change.

Changing the User's Password After Expiry

If your account has expired, you will be prompted with the following message "Your password expired. Please change your password by clicking here."

Click on the link and enter the login details with your expired password, to proceed to the Change Password page. For more information on setting passwords, Setting Password and Login Requirements , on page 263.

Note Changes to the password take effect immediately and do not require you to commit the change.

Managing Locally-Defined Administrative Users

Adding Locally-Defined Users

Follow this procedure to add users directly to the Security Management appliance if you are not using external authentication. Alternatively, use the **`userconfig`** command in the CLI.

Note If external authentication is also enabled, be sure that local user names do not duplicate externally-authenticated user names.

There is no limit to the number of user accounts that you can create on the appliance.

Step 1 If you will assign custom user roles, we recommend that you define those roles first. See Custom User Roles , on page 256.

Step 2 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 3 Choose **Management Appliance > System Administration > Users**.

Step 4 Click **Add User**.

Step 5 Enter a unique name for the user. You cannot enter words that are reserved by the system (such as "`operator`" and "`root`"). If you also use external authentication, user names should not duplicate externally-authenticated user names.

Step 6 Enter a full name for the user.

Step 7 Select a predefined role or a custom role. See the table *Descriptions of User Roles* in section Predefined User Roles , on page 254 for more information about user roles.
If you add a new Email role role here, enter a name for the role. For naming restrictions, see Creating Custom Email User Roles , on page 258.

Step 8 Enter a password and reenter it.

Step 9 Submit and commit your changes.

Step 10 If you added a custom user role on this page, assign privileges to that role now. See Custom User Roles , on page 256.

Editing Locally-Defined Users

Use this procedure to change a password, for example.

Step 1 Click the user's name in the Users listing.

Step 2 Make changes to the user.

Step 3 Submit and commit your changes.

Deleting Locally-Defined Users

Step 1 Click the trash can icon corresponding to the user's name in the Users listing.

Step 2 Confirm the deletion by clicking **Delete** in the warning dialog that appears.

Step 3 Click **Commit** to commit your changes.

Viewing the List of Locally-Defined Users

To view a list of locally-defined users, do the following:

- [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.
- Choose **Management Appliance > System Administration > Users**.

Note Asterisks indicate users assigned custom user roles for delegated administration. “Unassigned” appears in red if the user’s custom role has been deleted. For more information on custom user roles, see Custom User Roles , on page 256.

Setting and Changing Passwords

- When you add a user, you specify an initial password for that user.
- To change passwords for users configured on the system, use the Edit User page in the GUI (see Editing Locally-Defined Users, on page 262 for more information).
- To change the password for the default admin user account for the system, see Changing the Admin User's Password, on page 261.
- To force users to change their passwords, see Requiring Users to Change Passwords on Demand , on page 267.
- Users can change their own passwords by clicking the Options menu at the top right side of the GUI and selecting the Change Password option.

Setting Password and Login Requirements

You can define user account and password restrictions to enforce organizational password policies. The user account and password restrictions apply to local users defined on the Security Management appliance. You can configure the following settings:

- **User account locking.** You can define how many failed login attempts cause the user to be locked out of the account.
- **Password lifetime rules.** You can define how long a password can exist before the user is required to change the password after logging in.

- **Password rules.** You can define what kinds of passwords users can choose, such as which characters are optional or mandatory.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > Users**.

Step 3 Scroll down to the **Local User Account and Password Settings** section.

Step 4 Click **Edit Settings.**

Step 5 Configure settings:

Setting	Description
User Account Lock	Choose whether or not to lock the user account after the user fails to login successfully. Specify the number of failed login attempts that cause the account locking. You can enter any number from one (1) to 60. Default is five (5). When you configure account locking, enter the message to be displayed to the user attempting to login. Enter text using 7-bit ASCII characters. This message is only displayed when users enter the correct password to a locked account. When a user account gets locked, an administrator can unlock it on the Edit User page in the GUI or using the userconfig CLI command. Failed login attempts are tracked by user, regardless of the machine the user connects from or the type of connection, such as SSH or HTTP. Once the user successfully logs in, the number of failed login attempts is reset to zero (0). When a user account is locked out due to reaching the maximum number of failed login attempts, an alert is sent to the administrator. The alert is set at the “Info” severity level. **Note** You can also manually lock individual user accounts. See Locking User Accounts Manually, on page 267.
Password Reset	Choose whether or not users should be forced to change their passwords after an administrator changes their passwords. You can also choose whether or not users should be forced to change their passwords after they expire. Enter the number of days a password can last before users must change it. You can enter any number from one (1) to 366. Default is 90. To force users to change their passwords at non-scheduled times, see Requiring Users to Change Passwords on Demand , on page 267. When you force users to change their passwords after they expire, you can display a notification about the upcoming password expiration. Choose the number of days before expiration to notify users. **Note** When a user account uses SSH keys instead of a password challenge, the Password Reset rules still apply. When a user account with SSH keys expires, the user must enter their old password or ask an administrator to manually change the password to change the keys associated with the account.

Setting	Description
Password Rules: Require at least <number> characters.	Enter the minimum number of characters that passwords may contain. Enter any number between zero (0) and 128. The default is 8. Passwords can have more characters than the number you specify here.
Password Rules: Require at least one number (0-9).	Choose whether or not the passwords must contain at least one number.
Password Rules: Require at least one special character.	Choose whether or not the passwords must contain at least one special character. Passwords may contain the following special characters: ~ ? ! @ # $ % ^ & * - _ + = \ | / [ ] ( ) < > { } ` ' " ; : , .
Password Rules: Ban usernames and their variations as passwords.	Choose whether or not the password are allowed to be the same as the associated user name or variations on the user name. When user name variations are banned, the following rules apply to passwords: • The password may not be the same as the user name, regardless of case. • The password may not be the same as the user name in reverse, regardless of case. • The password may not be the same as the user name or reversed user name with the following character substitutions: ◦ "@" or "4" for "a" ◦ "3" for "e" ◦ "|", "!", or "1" for "i" ◦ "0" for "o" ◦ "$" or "5" for "s" ◦ "+" or "7" for "t"
Password Rules: Ban reuse of the last <number> passwords.	Choose whether or not users are allowed to choose a recently used password when they are forced to change the password. If they are not allowed to reuse recent passwords, enter the number of recent passwords that are banned from reuse. You can enter any number from one (1) to 15. Default is three (3).

Setting	Description
Password Rules: List of words to disallow in passwords	You can create a list of words to disallow in passwords. Make this file a text file with each forbidden word on a separate line. Save the file with the name forbidden_password_words.txt and use SCP or FTP to upload the file to the appliance. If this restriction is selected but no word list is uploaded, this restriction is ignored.
Password Strength	You can display a password-strength indicator when an admin or user enters a new password. This setting does not enforce creation of strong passwords, it merely shows how easy it is to guess the entered password. Select the roles for which you wish to display the indicator. Then, for each selected role, enter a number greater than zero. A larger number means that a password that registers as strong is more difficult to achieve. This setting has no maximum value. Examples: • If you enter 30 , then an 8 character password with at least one upper- and lower-case letter, number, and special character will register as a strong password. • If you enter 18 , then an 8 character password with all lower case letters and no numbers or special characters will register as strong. Password strength is measured on a logarithmic scale. Evaluation is based on the U.S. National Institute of Standards and Technology rules of entropy as defined in NIST SP 800-63, Appendix A. Generally, stronger passwords: • Are longer • Include upper case, lower case, numeric, and special characters • Do not include words in any dictionary in any language. To enforce passwords with these characteristics, use the other settings on this page.

Step 6 Submit and commit your changes.

What to Do Next

Require users to change their passwords to new passwords that meet the new requirements. See Requiring Users to Change Passwords on Demand , on page 267

Requiring Users to Change Passwords on Demand

To require all or selected users to change their passwords at any time on an ad-hoc basis, perform the steps in this procedure. This is a one-time action.

To automate a periodic requirement for changing passwords, use the Password Reset option described in Setting Password and Login Requirements , on page 263.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > Users**.

Step 3 In the Users section, select the check boxes beside the users who will be required to change passwords .

Step 4 Select **Enforce Password Changes**.

Step 5 Select options.
The global setting for the grace period is configured in Local User Account & Password Settings.

Step 6 Click **OK**.

Locking and Unlocking Local User Accounts

Locking a user account prevents a local user from logging into the appliance. A user account can be locked in one of the following ways:

- You can configure all local user accounts to lock after users fail to log in successfully after a configured number of attempts. See Setting Password and Login Requirements , on page 263.
- Administrators can manually lock user accounts. See Locking User Accounts Manually, on page 267.

AsyncOS displays the reason why the user account was locked when you view the user account on the Edit User page.

Locking User Accounts Manually

Step 1 First time only: Set up the appliance to enable user account locking:

Step 2
a) [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.
b) Go to **Management Appliance > System Administration > Users**.
c) In the **Local User Account & Password Settings** section, click **Edit Settings**.
d) Select the checkbox to **Display Locked Account Message if Administrator has manually locked a user account** and enter your message.
e) Submit the change.

Step 3 Go to **Management Appliance > System Administration > Users** and click the user name.

Note Before you lock the Admin account, be sure that you can unlock it. See the Note in Unlocking User Accounts , on page 268.

Step 4 Click **Lock Account**.
AsyncOS displays a message saying that the user will be unable to log into the appliance and asks if you want to continue.

Unlocking User Accounts

To unlock a user account, open the user account by clicking on the user name in the Users listing and click Unlock Account.

Note If you lock the admin account, you can only unlock it by logging in as the admin through a serial communications connection to the serial console port. The admin user can always access the appliance using the serial console port, even when the admin account is locked. See the "Setup and Installation" chapter in the documentation or online help for your Email Security appliance for more information on accessing the appliance using the serial console port.

External User Authentication

If you store user information in an LDAP or RADIUS directory on your network, you can configure your Security Management appliance to use the external directory to authenticate users who log in to the appliance.

Note Some features described in Customizing Your View , on page 345 are not available to externally-authenticated users.

- If your deployment uses both local and external authentication, local user names must not duplicate externally-authenticated user names.
- If the appliance cannot communicate with the external directory, a user who has both an external and a local account can log in with a local user account on the appliance.

See:

- Configuring External Authentication of Administrative Users Using LDAP , on page 242
- Enabling RADIUS Authentication, on page 268

Configuring LDAP Authentication

To configure LDAP authentication, see Configuring External Authentication of Administrative Users Using LDAP , on page 242.

Enabling RADIUS Authentication

You can use a RADIUS directory to authenticate users and assign groups of users to user roles for administering your appliance. The RADIUS server should support the CLASS attribute, which AsyncOS uses to assign users in the RADIUS directory to user roles.

Note If an external user changes the user role for their RADIUS group, the user should log out of the appliance and then log back in. The user will have the permissions of their new role.

Before You Begin

The Shared Secret key for access to the RADIUS server must be no more than 48 characters long.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance** > **System Administration** > **Users** page and click **Enable**.

Step 3 Select the **Enable External Authentication** check box.

Step 4 Select RADIUS for the authentication type.

Step 5 Enter the host name for the RADIUS server.

Step 6 Enter the port number for the RADIUS server. The default port number is `1812`.

Step 7 Enter the Shared Secret key for the RADIUS server.

Note When enabling external authentication for a cluster of Email Security appliances, enter the same Shared Secret key on all appliances in the cluster.

Step 8 Enter the number of seconds that the appliance waits for a response from the server before timing out.

Step 9 Select whether to use Password Authentication Protocol (PAP) or Challenge Handshake Authentication Protocol (CHAP) for the authentication protocol.

Step 10 (Optional) Click **Add Row** to add another RADIUS server. Repeat steps 6 and 7 for each RADIUS server that your appliance uses for authentication.
When you define multiple external servers, the appliance connects to the servers in the order defined on the appliance. You might want to define multiple external servers to allow for failover in case one server is temporarily unavailable.

Step 11 Enter the amount of time to store external authentication credentials in the web user interface.

Note If the RADIUS server uses one-time passwords, for example passwords created from a token, enter zero (0). When the value is set to zero, AsyncOS does not contact the RADIUS server again to authenticate during the current session.

Step 12 Configure Group Mapping:

Setting	Description
Map externally authenticated users to multiple local roles (Recommended)	AsyncOS assigns RADIUS users to appliance roles based on the RADIUS CLASS attribute. CLASS attribute requirements: • 3 character minimum • 253 character maximum • no colons, commas, or newline characters • one or more mapped CLASS attributes for each RADIUS user (With this setting, AsyncOS denies access to RADIUS users without a mapped CLASS attribute.) For RADIUS users with multiple CLASS attributes, AsyncOS assigns the most restrictive role. For example, if a RADIUS user has two CLASS attributes, which are mapped to the Operator and Read-Only Operator roles, AsyncOS assigns the RADIUS user to the Read-Only Operator role, which is more restrictive than the Operator role. These are the appliance roles ordered from least restrictive to most restrictive: • Administrator • Email Administrator • Custom user role (email) If a user is assigned multiple Class attributes that are mapped to custom user roles, the last class attribute on the list on the RADIUS server will be used. • Technician • Operator • Read-Only Operator • Help Desk User • Guest
Map all externally authenticated users to the Administrator role	AsyncOS assigns RADIUS users to the Administrator role.

Step 13 (Optional) Click **Add Row** to add another group. Repeat step 11 for each group of users that the appliance authenticates.

Step 14 Submit and commit your changes.

Two-Factor Authentication

You can use a RADIUS directory to configure two-factor authentication for specific user roles. The appliance supports the following authentication protocols for communicating with the RADIUS server:

- Password Authentication Protocol (PAP)
- Challenge Handshake Authentication Protocol (CHAP)

You can enable two-factor authentication for the following user roles:

- predefined
- custom

The feature has been tested with:

- RSA Authentication Manager v8.2
- FreeRADIUS v1.1.7 and above
- ISE v1.4 and above

Related Topics:

Enabling Two-Factor Authentication

Make sure that you get the required RADIUS server details for two-factor authentication from your IT administrator.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **System Administration > Users** page and click **Enable** under Two-Factor Authentication.

Step 3 Enter the hostname or IP address of the RADIUS server.

Step 4 Enter the port number of the RADIUS server.

Step 5 Enter the Shared Secret password of the RADIUS server.

Step 6 Enter the number of seconds to wait for a response from the server before timing out.

Step 7 Select the appropriate authentication protocol.

Step 8 (Optional) Click **Add Row** to add another RADIUS server. Repeat steps 2 to 6 for each RADIUS server.

Note You can add up to ten RADIUS servers.

Step 9 Select the required user roles for which you want to enable two-factor authentication.

Step 10 Submit and commit your changes.

When two-factor authentication is enabled, the user is prompted to enter a passcode after entering the username and passphrase, to login to the appliance.

Disabling Two-Factor Authentication

Before You Begin

Make sure you have enabled two-factor authentication on your appliance.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **System Administration > Users** page and click **Edit Global Settings** under Two-Factor Authentication

Step 3 Deselect **Enable Two-Factor Authentication**.

Step 4 Submit and commit your changes.

Adding an Email or Web Security appliance over SSH with Pre-Shared Keys

The following example demonstrates how to add an Email Security applaince (testesa.example.com) to the Security Management appliance (testsma.example.com) over SSH using pre-shared keys.

To add a Web Security appliance, choose **WSA** when prompted to enter the type of Cisco appliance.

```
testsma.example.com> applianceconfig

Choose the operation you want to perform.

ADD - Add SMA Connection Parameters and Keys.
EDIT - Edit an appliance.
DELETE - Remove an appliance.
TEST - Test that an appliance is properly configured.
SERVICES - Configure the centralized services for an appliance.
STATUS - Display the status of centralized services.
PORT - Configure which port is used to communicate with remote appliances.

[]> add

Please enter the type of Cisco appliance that this device is
1. ESA
2. WSA

[1]> 1

Enter the IP address or hostname of an appliance to transfer data with.
(A hostname may be entered in this field, however it will be immediately
resolved to an IP address when the form is submitted.)
[]> IP address entered

Enter a name to identify this appliance

[]> name of appliance

File transfer access via SSH is required to transfer reporting data, message logs,
and quarantine safelist/blocklist data from appliances

Would you like to configure file transfer access for this appliance? [Y]>

Would you like to use a custom ssh port to connect to this appliance? [N]>

Would you like to connect an Email Security appliance using pre-shared keys?
Use this option if you have enabled two-factor authentication on the Email
Security appliance. [N]> yes
```

```
To add an Email Security appliance to the Content Security Management appliance
using pre-shared keys, log in to the Email Security appliance,
run the smaconfig > add command, enter the following details.

Host: vm10sma0006.qa

User Key:
AAAAB3NzaC1yc2EAAAADAQABAAABAQDgm3kG9RHc4gVZxRe0orh5DW5Yje5UB9BpJqcTRQJoxUIAv2Xig
8q5geyaWHZcFoUxH61YQbPX3R8CVMYgJ8/QB/iunjkr3jowV/SCuBBikEFgj1zuxlsFhL0L487epEgby1gH0rfJ
gwSa2/6dhfyUayst6pT87CZGOQltgx7s51wc+ve770X3SqlQD5bdYC4x9+gCX0wdwfhTH1+4/82jwYjK1lAEXc
O4k4TuZJEJnyBQ3YyCyVwXuDkXpI6xJDemxcc36e7Wwtpn3mn2VLaTG2/I38XwSvlYB6TcqmWnO10gL+aD
wkKAKcuhYpz4NFr9myej1mhMk7ZAFXmRNxvT
```

Note Before you proceed to the next step, ensure that you have added the **Host** and **User Key** details to the Email or Web Security appliance. Commit the changes in your Email or Web Security appliance before continuing the process of adding connection parameters in the Security Management appliance.

```
Do you want to continue connecting using pre-shared keys? [Y]> yes
```

Additional Controls on Access to the Security Management Appliance

Configuring IP-Based Network Access

You can control from which IP addresses users access the Security Management appliance by creating access lists for users who connect directly to the appliance and users who connect through a reverse proxy, if your organization uses reverse proxies for remote users.

Direct Connections

You can specify the IP addresses, subnets, or CIDR addresses for machines that can connect to the Security Management appliance. Users can access the appliance from any machine with IP address from the access list. Users attempting to connect to the appliance from an address not included in the list are denied access.

Connecting Through a Proxy

If your organization's network uses reverse proxy servers between remote users' machines and the Security Management appliance, AsyncOS allows you create an access list with the IP addresses of the proxies that can connect to the appliance.

Even when using a reverse proxy, AsyncOS still validates the IP address of the remote user's machine against a list of IP addresses allowed for user connections. To send the remote user's IP address to the Email Security appliance, the proxy needs to include the x-forwarded-for HTTP header in its connection request to the appliance.

The x-forwarded-for header is a non-RFC standard HTTP header with the following format:

x-forwarded-for: client-ip, proxy1, proxy2,... CRLF .

The value for this header is a comma-separated list of IP addresses with the left-most address being the address of the remote user's machine, followed by the addresses of each successive proxy that forwarded the connection

request. (The header name is configurable.) The Security Management appliance matches the remote user's IP address from the header and the connecting proxy's IP address against the allowed user and proxy IP addresses in the access list.

Note AsyncOS supports only IPv4 addresses in the x-forwarded-for header.

Creating the Access List

You can create the network access list either via the Network Access page in the GUI or the **`adminaccessconfig > ipaccess`** CLI command. The following figure shows the Network Access page with a list of user IP addresses that are allowed to connect directly to the Security Management appliance.

Figure 11: Example Network Access Settings

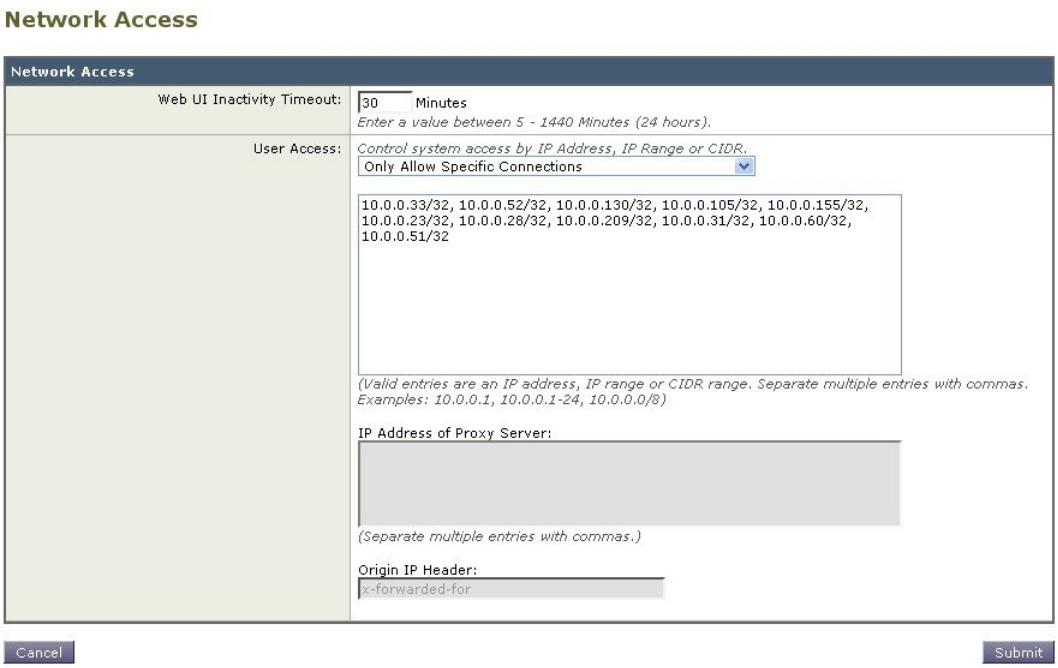

AsyncOS offers four different modes of control for the access list:

- **Allow All.** This mode allows all connections to the appliance. This is the default mode of operation.
- **Only Allow Specific Connections.** This mode allows a user to connection to the appliance if the user's IP address matches the IP addresses, IP ranges, or CIDR ranges included in the access list.
- **Only Allow Specific Connections Through Proxy.** This mode allows a user to connect to the appliance through a reverse proxy if the following conditions are met:
  - The connecting proxy's IP address is included in the access list's IP Address of Proxy Server field.
  - The proxy includes the x-forwarded-header HTTP header in its connection request.
  - The value of x-forwarded-header is not empty.

- The remote user's IP address is included in x-forwarded-header and it matches the IP addresses, IP ranges, or CIDR ranges defined for users in the access list.

- **Only Allow Specific Connections Directly or Through Proxy.** This mode allows users to connect through a reverse proxy or directly to the appliance if their IP address matches the IP addresses, IP ranges, or CIDR ranges included in the access list. The conditions for connecting through a proxy are the same as in the Only Allow Specific Connections Through Proxy mode.

Please be aware that you may lose access to the appliance after submitting and committing your changes if one of the following conditions is true:

- If you select **Only Allow Specific Connections** and do not include the IP address of your current machine in the list.
- If you select **Only Allow Specific Connections Through Proxy** and the IP address of the proxy currently connected to the appliance is not in the proxy list and the value of the Origin IP header is not in the list of allowed IP addresses.
- If you select **Only Allow Specific Connections Directly or Through Proxy** and
  - the value of the Origin IP header is not in the list of allowed IP addresses

    OR
  - the value of the Origin IP header is not in the list of allowed IP Addresses and the IP address of the proxy connected to the appliance is not in the list of allowed proxies.

If you choose to continue without correcting the access list, AsyncOS will disconnect your machine or proxy from the appliance when you commit your changes.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **System Administration > Network Access**.

Step 3 Click **Edit Settings**.

Step 4 Select the mode of control for the access list.

Step 5 Enter the IP addresses from which users will be allowed to connect to the appliance.
You can enter an IP address, IP address range or CIDR range. Use commas to separate multiple entries.

Step 6 If connecting through a proxy is allowed, enter the following information:

- The IP addresses of the proxies allowed to connect to the appliance. Use commas to separate multiple entries.
- The name of the origin IP header that the proxy sends to the appliance, which contains the IP addresses of the remote user's machine and the proxy servers that forwarded the request. By default, the name of the header is x-forwarded-for .

Step 7 Submit and commit your changes.

Configuring the Web UI Session Timeout

You can specify how long a user can be logged into the Security Management appliance's Web UI before AsyncOS logs the user out due to inactivity. This Web UI session timeout applies to all users, including admin, and it is used for both HTTP and HTTPS sessions.

Once AsyncOS logs a user out, the appliance redirects the user's web browser to login page.

Note The Web UI Session Timeout does not apply to spam quarantine sessions, which have a 30 minute timeout that cannot be configured.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Use the **System Administration > Network Access** page.

Step 3 Click **Edit Settings**.

Step 4 Enter the number of minutes users can be inactive before being logged out. You can define a timeout period between 5 and 1440 minutes.

Step 5 Submit and commit your changes.

Controlling Access to Sensitive Information in Message Tracking

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Go to the **Management Appliance > System Administration > Users** page.

Step 3 In the **Tracking Privileges** section, click **Edit Settings**.

Step 4 Select the roles for which you want to grant access to sensitive information in Message Tracking.
Only custom roles with access to Message Tracking are listed.

Step 5 Submit and commit your changes.
The Centralized Email Message Tracking feature must be enabled under Management Appliance > Centralized Services for this setting to take effect.

Displaying a Message for Administrative Users

You can display a message that administrative users will see when they sign in to the appliance.

To set or clear a message:

Step 1 If you will import a text file, put it into the `/data/pub/configuration` directory on the appliance.
Step 2 Access the command-line interface (CLI).
Step 3 Use the `adminaccessconfig > BANNER` command and subcommand.
Step 4 Commit the change.

Viewing Administrative User Activity

Viewing Active Sessions Using the Web

From the Security Management appliance, you can view all active sessions and users logged in to the appliance.

From the upper right corner of the window, choose **Options > Active Sessions**.

Figure 12: Active Sessions Menu

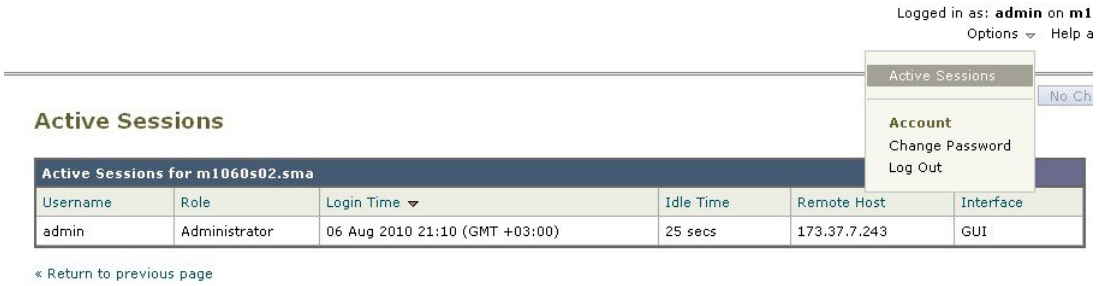

From the Active Sessions page you can view the User name, what role the user has, the time the user logged in, idle time, and whether the user is logging in from the command line or the GUI.

Viewing Your Recent Login Attempts

To view your last few recent login attempts (failed or successful) via the web interface, SSH, and/or FTP:

Step 1 Log in.
Step 2 Click the Figure-icon icon beside "Logged in as" near the top right side of the screen.

Viewing Administrative User Activity via the Command Line Interface

The following commands support multiuser access to the appliance.

- The **who** command lists all users who are logged in to the system via the CLI or the web user interface, the role of the user, the time of login, the idle time, and the remote host from which the user is logged in.
- The whoami command displays the user name and full name of the user currently logged in, and which groups the user belongs to:

```
mail3.example.com>
whoami
Username: admin
Full Name: Administrator
Groups: admin, operators, config, log, guest
```

- The last command displays which users have recently logged into the appliance. The IP address of the remote host, and the login, logout, and total time also appear.

```
mail3.example.com> last
Username  Remote Host  Login Time         Logout Time        Total Time
========  ===========  ================   ================   ==========
admin     10.1.3.67    Sat May 15 23:42   still logged in    15m
admin     10.1.3.67    Sat May 15 22:52   Sat May 15 23:42   50m
admin     10.1.3.67    Sat May 15 11:02   Sat May 15 14:14   3h 12m
admin     10.1.3.67    Fri May 14 16:29   Fri May 14 17:43   1h 13m
shutdown                                  Fri May 14 16:22
shutdown                                  Fri May 14 16:15
admin     10.1.3.67    Fri May 14 16:05   Fri May 14 16:15   9m
admin     10.1.3.103   Fri May 14 16:12   Fri May 14 16:15   2m
admin     10.1.3.103   Thu May 13 09:31   Fri May 14 14:11   1d 4h 39m
admin     10.1.3.135   Fri May 14 10:57   Fri May 14 10:58   0m
admin     10.1.3.67    Thu May 13 17:00   Thu May 13 19:24   2h 24m
```

Troubleshooting Administrative User Access

Error: User Has No Access Privileges Assigned

Problem

A user to whom you have delegated administration can log in to the Security Management appliance but sees a message that no access privileges are assigned.

Solution

Make sure that you have assigned privileges to the custom user role to which this user is assigned. Look at Management Appliance > System Administration > Users to determine the User Role assigned, then go to Management Appliance > System Administration > User Roles, click the name of the User Role, and assign privileges to the role.

If you have assigned access based on Reporting Group, make sure you have selected a Reporting Group for that user on the Management Appliance > System Administration > User Roles page. To assign a group, click the **No groups selected** link in the Email Reporting column of the User Roles for Delegated Administration table.

User Has No Active Menus

Problem

A user to whom you have granted Publish privileges has no active menus upon login.

Solution

Make sure you have granted access to at least one Access Policy or Custom URL Category. If you do not want to grant this user privileges to edit either of these, create a custom URL category which is not used in any policy and grant this user role privileges to this category on the Custom User Role page.

Externally-Authenticated Users See Preferences Option

Problem

Externally-authenticated users see the Preferences option.

Solution

Ensure that users that you add directly in the Security Management appliance have unique usernames that are not also used in your external authentication database.

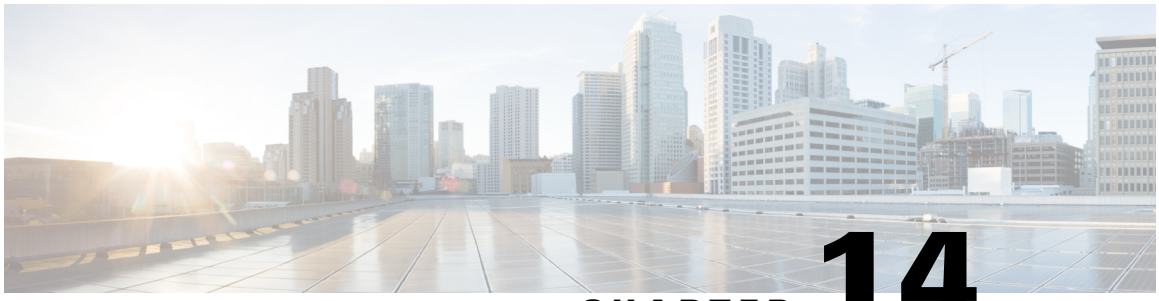

CHAPTER 14

Common Administrative Tasks

This chapter contains the following sections:

Performing Administrative Tasks

You can perform most system administration tasks by using the System Administration menu in the graphical user interface (GUI). Some system administration features, however, are available only in the command-line interface (CLI).

In addition, you access the status-monitoring features of the appliance on the Monitor menu, which is described in chapter Monitoring System Status, on page 223

Note Several of the features or commands described in this chapter can affect routing precedence. For more information, see IP Addresses, Interfaces, and Routing, on page 399.

Working with Feature Keys

Keys are specific to the serial number of your appliance and specific to the feature that you enable. You cannot reuse a key from one system on another system.

To perform the tasks described in this section from the command-line prompt, use the featurekey command.

To	Do This
• View all active feature keys for the appliance • View any feature keys that are pending activation • Search for new keys that have been issued • Install feature keys manually • Activate feature keys	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface. Select **Management Appliance > System Administration > Feature Keys**. To add a new feature key manually, paste or enter the key into the Feature Key field and click **Submit Key.** An error message appears if the feature is not added (for example, if the key is incorrect); otherwise, the feature key is added to the list. If the appliance is configured to automatically download and install new keys as they are issued. the Pending Activation list is always empty.
Enable or disable automatic download and activation of feature keys	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface. Select **Management Appliance > System Administration > Feature Keys Settings** By default, the appliance periodically checks for new keys.
Renew expired feature keys	Contact your Cisco representative

Virtual Appliance Licensing and Feature Keys

For information about appliance behavior upon license and feature key expiration, see the *Cisco Content Security Virtual Appliance Installation Guide* available from http://www.cisco.com/c/en/us/support/security/content-security-management-appliance/products-installation-guides-list.html

To view license information, use the show license command in the command-line interface (CLI.)

Performing Maintenance Tasks Using CLI Commands

The operations and commands described in this section enable you to perform maintenance-related tasks on the Security Management appliance. This section describes the following operations and commands:

- shutdown
- reboot
- suspend
- suspendtransfers
- resume
- resumetransfers
- resetconfig
- version

Shutting Down the Security Management Appliance

To shut down your Security Management appliance, do the folliwng:

- [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.
- Use the **Management Appliance > System Administration > Shutdown/Reboot** page.

  or

- Use the `shutdown` command at the command-line prompt.

Shutting down an appliance exits AsyncOS, which allows you to safely power down the appliance. You may restart the appliance at a later time without losing any messages in the delivery queue. You must enter a delay for the appliance to shut down. The default delay is 30 seconds. AsyncOS allows open connections to complete during the delay, after which it forcefully closes open connections.

Rebooting the Security Management Appliance

To reboot your Security Management appliance, use the Shutdown/Reboot page available on the System Administration menu in the GUI, or use the r eboot command in the CLI.

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

Taking the Security Management Appliance Out of Service

If you want to take the appliance offline, for example to perform system maintenance, use one of the following commands:

Command	Description	Persistence
`suspend`	• Suspends transfer of quarantined messages from the Email Security appliance to the Security Management appliance. • Suspends delivery of messages released from quarantines. • Inbound email connections are not accepted. • Outbound email delivery is halted. • Log transfers are halted. • The CLI remains accessible.	Persists after reboot.
`suspendtransfers`	Suspends transfer of reporting and tracking data from managed email security appliances to the content security management appliance. This command also suspends receiving of quarantined messages from Email Security appliances. Use this command when preparing to bring a backup appliance into service as the primary appliance.	Persists after reboot.

You must enter a delay for the appliance when using these commands. The default delay is 30 seconds. AsyncOS allows open connections to complete during the delay, after which it forcefully closes open connections. If there are no open connections, service is suspended immediately.

To re-activate services that were halted by the suspend or suspendtransfers commands, use the resume or resumetransfers commands, respectively.

To determine the current online/suspended status of the management appliance, select **Management Appliance > System Administration > Shutdown/Reboot** in the web interface.

See also:

- "Suspending Email Delivery," "Resuming Email Delivery," "Suspending Receiving," and "Resuming Receiving" in the documentation or online help for your Email Security appliance.

CLI Examples: suspend and suspendtransfers Commands

```
sma.example.com> suspend
Enter the number of seconds to wait before abruptly closing connections.
[30]> 45
Waiting for listeners to exit...
Receiving suspended.
Waiting for outgoing deliveries to finish...
Mail delivery suspended.
sma.example.com>
sma.example.com> suspendtransfers

Transfers suspended.
sma.example.com>
```

Resuming from a Suspended State

The resume command returns the appliance to normal operating state after using the suspend or suspenddel command.

The resumetransfers command returns the appliance to normal operating state after using the suspendtransfers command.

CLI Examples: resume and resumetransfers Commands

```
sma.example.com> resume
Receiving resumed.
Mail delivery resumed.
sma.example.com>
sma.example.com> resumetransfers

Receiving resumed.
Transfers resumed.
sma.example.com>
```

Resetting the Configuration to Factory Defaults

When physically transferring the appliance, or as a last resort for solving configuration issues, you may want to reset the appliance to factory defaults.

Caution Resetting the configuration will disconnect you from the CLI, disable services that you used to connect to the appliance (FTP, Telnet, SSH, HTTP, HTTPS), and remove user accounts.

To	Do This
• Reset all configurations to factory defaults • Clear all reporting counters But • Retain log files • Retain quarantined messages	1 Ensure that you can connect to the appliance after reset using the default admin user account and password, either to the CLI using the serial interface or to the Management port using the default settings. See chapter Setup, Installation, and Basic Configuration, on page 9 for information about accessing an appliance having default settings. 2 Suspend service on the appliance. 3 Select **Management Appliance > System Administration > Configuration File** and click **Reset**. **Note** After resetting, the appliance automatically returns to the online state. If mail delivery was suspended before reset, delivery will be attempted again after the reset.
• Reset all configurations to factory defaults • Remove all data	Use the diagnostic > reload CLI command. **Caution** This command is NOT the same as the similar command used on a Cisco router or switch.

The resetconfig Command

```
mail3.example.com> suspend
Delay (seconds, minimum 30):
[30]> 45
Waiting for listeners to exit...
Receiving suspended.
Waiting for outgoing deliveries to finish...
Mail delivery suspended.
mail3.example.com> resetconfig
Are you sure you want to reset all configuration values? [N]> Y
All settings have been restored to the factory default.
```

Displaying the Version Information for AsyncOS

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliances > Centralized Services > System Status**.

Step 3 Scroll to the bottom of the page and look under Version Information to see the version of AsyncOS that is currently installed.
Additionally, you can use the **version** command at the command-line prompt.

Enabling Remote Power Cycling

The ability to remotely reset the power for the appliance chassis is available only on 80- and 90- series hardware.

If you want to be able to remotely reset appliance power, you must enable and configure this functionality in advance, using the procedure described in this section.

Before You Begin

- Cable the dedicated Remote Power Cycle (RPC) port directly to a secure network. For information, see hardware documentation for your model, available from the location listed in Documentation , on page 405.
- Ensure that the appliance is accessible remotely; for example, open any necessary ports through the firewall.
- This feature requires a unique IPv4 address for the dedicated Remote Power Cycle interface. This interface is configurable only via the procedure described in this section; it cannot be configured using the ipconfig command.
- In order to cycle appliance power, you will need a third-party tool that can manage devices that support the Intelligent Platform Management Interface (IPMI) version 2.0. Ensure that you are prepared to use such a tool.
- For more information about accessing the command-line interface, see the CLI reference guide.

Step 1 Use SSH, telnet, or the serial console port to access the command-line interface.

Step 2 Sign in using an account with Administrator access.

Step 3 Enter the following commands:
remotepower

setup

Step 4 Follow the prompts to specify the following:

- The dedicated IP address for this feature, plus netmask and gateway.
- The username and password required to execute the power-cycle command.

These credentials are independent of other credentials used to access your appliance.

Step 5 Enter commit to save your changes.

Step 6 Test your configuration to be sure that you can remotely manage appliance power.

Step 7 Ensure that the credentials that you entered will be available to you in the indefinite future. For example, store this information in a safe place and ensure that administrators who may need to perform this task have access to the required credentials.

What to Do Next

Remotely Resetting Appliance Power , on page 388

Monitoring System Health Using SNMP

AsyncOS supports system status monitoring via Simple Network Management Protocol (SNMP) versions v1, v2, and v3.

- To enable and configure SNMP, use the snmpconfig command in the command-line interface.
- MIBs are available from http://www.cisco.com/c/en/us/support/security/content-security-management-appliance/tsd-products-support-series-home.html use the latest available files.
- The use of SNMPv3 with password authentication and DES Encryption is mandatory to enable this service. (For more information on SNMPv3, see RFCs 2571-2575.) You are required to set a SNMPv3 passphrase of at least 8 characters to enable SNMP system status monitoring. The first time you enter a SNMPv3 passphrase, you must re-enter it to confirm. The snmpconfig command “remembers” this phrase the next time you run the command.
- When setting up SNMP to monitor connectivity:

  When entering the url-attribute while configuring a connectivityFailure SNMP trap, determine whether the URL is pointing at a directory or a file.

  - If it is a directory, add a trailing slash (/)
  - If it is a file, do not add a trailing slash

- Additional information about using SNMP with AsyncOS is available in the online help for your email security appliance.

Example: snmpconfig Command

```
sma.example.com> snmpconfig
Current SNMP settings:
SNMP Disabled.
Choose the operation you want to perform:
- SETUP - Configure SNMP.
[]> SETUP
Do you want to enable SNMP?
[Y]>
Please choose an IP interface for SNMP requests.
1. Management (198.51.100.1: sma.example.com)
[1]>
Which port shall the SNMP daemon listen on interface "Management"?
[161]>
Please select SNMPv3 authentication type:
1. MD5
2. SHA
[1]> 2
Please select SNMPv3 privacy protocol:
1. DES
2. AES
[1]> 2
Enter the SNMPv3 authentication passphrase.
[]>
Please enter the SNMPv3 authentication passphrase again to confirm.
```

```
[]>
Enter the SNMPv3 privacy passphrase.
[]>
Please enter the SNMPv3 privacy passphrase again to confirm.
[]>
Service SNMP V1/V2c requests?
[N]> Y
Enter the SNMP V1/V2c community string.
[ironport]> public
Shall SNMP V2c requests be serviced from IPv4 addresses?
[Y]>
From which IPv4 networks shall SNMP V1/V2c requests be allowed? Separate
multiple networks with commas.
[127.0.0.1/32]>
Enter the Trap target as a host name, IP address or list of IP
addresses separated by commas (IP address preferred). Enter "None" to disable traps.
[127.0.0.1]> 203.0.113.1
Enter the Trap Community string.
[ironport]> tcomm
Enterprise Trap Status
1. CPUUtilizationExceeded       Disabled
2. FIPSModeDisableFailure       Enabled
3. FIPSModeEnableFailure        Enabled
4. FailoverHealthy              Enabled
5. FailoverUnhealthy            Enabled
6. RAIDStatusChange             Enabled
7. connectivityFailure          Disabled
8. fanFailure                   Enabled
9. highTemperature              Enabled
10. keyExpiration               Enabled
11. linkUpDown                  Enabled
12. memoryUtilizationExceeded   Disabled
13. powerSupplyStatusChange     Enabled
14. resourceConservationMode    Enabled
15. updateFailure               Enabled
Do you want to change any of these settings?
[N]> Y
Do you want to disable any of these traps?
[Y]> n
Do you want to enable any of these traps?
[Y]> y
Enter number or numbers of traps to enable.  Separate multiple numbers with
commas.
[]> 1,7,12
What threshold would you like to set for CPU utilization?
[95]>
What URL would you like to check for connectivity failure?
[http://downloads.ironport.com]>
What threshold would you like to set for memory utilization?
[95]>
Enter the System Location string.
[Unknown: Not Yet Configured]> Network Operations Center - west; rack #30, position 3
Enter the System Contact string.
[snmp@localhost]> SMA.Administrator@example.com
Current SNMP settings:
Listening on interface "Management" 198.51.100.1  port 161.
SNMP v3: Enabled.
SNMP v1/v2: Enabled, accepting requests from subnet 127.0.0.1/32 .
SNMP v1/v2 Community String: public
Trap target: 203.0.113.1
Location: Network Operations Center - west; rack #30, position 3
System Contact: SMA.Administrator@example.com
Choose the operation you want to perform:
- SETUP - Configure SNMP.
[]>
sma.example.com> commit
Please enter some comments describing your changes:
[]> Enable and configure SNMP
Changes committed: Fri Nov 06 18:13:16 2015 GMT
sma.example.com>
```

Backing Up Security Management Appliance Data

What Data Is Backed Up

You can choose to back up all data, or any combination of the following data:

- Spam quarantine, including messages and meta data
- Centralized policy, virus, and outbreak quarantines, including messages and meta data
- Email tracking (message tracking), including messages and meta data
- Reporting (Email)
- Safelist/blocklist

After the data transfer is finished, the data on the two appliances will be identical.

Configurations and logs are not backed up using this process. To back up those items, see Other Important Backup Tasks , on page 295.

Each backup after the first backup copies only the information generated since the last backup.

Restrictions and Requirements for Backups

Be sure to address the following restrictions and requirements before you schedule a backup:

Restriction	Requirement
AsyncOS version	The AsyncOS version of the source and target Security Management appliances must be the same. If there is a version incompatibility, upgrade appliances to the same release before scheduling a backup.
Target appliance on the network	The target appliance must be set up on the network. If the target appliance is new, run the System Setup Wizard to enter the necessary information. For instructions, see Setup, Installation, and Basic Configuration, on page 9
Communication between source and target appliances	The source and target Security Management appliances must be able to communicate using SSH. Therefore: • Port 22 must be open on both appliances. By default, this port is opened when you run the System Setup Wizard. • The Domain Name Server (DNS) must be able to resolve the host names of both appliances using both A records and PTR records.

Restriction	Requirement
Target appliance must not be in service	Only the primary appliance should pull data from managed email security appliances. To ensure this, see Prevent the Target Appliance From Pulling Data Directly from Managed Appliances , on page 293. Also, cancel any scheduled configuration publishing jobs on the backup appliance.
Appliance capacity	The disk space capacity of the target appliance must be the same as or greater than the capacity of the source appliance. Disk space allocated to each type of data (reporting, tracking, quarantine, etc.) on the target appliance cannot be less than the corresponding allocation on the source appliance. You can schedule a backup from a larger source to a smaller target Security Management appliance as long as there is enough space on the target appliance for all of the data being backed up, for each type of data. If the source appliance is larger than the target appliance, you must reduce the space allocated on the source appliance to match the space available on the smaller target appliance. To view and manage disk space allocations and capacity, see Managing Disk Space , on page 333. For disk capacity of virtual appliances, see the *Cisco Content Security Virtual Appliance Installation Guide* .
Multiple, concurrent, and chained backups	Only one backup process can run at a time; a backup that is scheduled to run before a previous backup has been completed will be skipped and a warning sent. Data from a Security Management appliance can be backed up to a single Security Management appliance. Chained backup (a backup to a backup) is not supported.

Backup Duration

During a full initial backup, a backup of 800GB may take up to 10 hours. Daily backups, may take up to 3 hours each. Weekly and monthly backups may take longer. These numbers may vary.

After the initial backup, the backup process transfers only files that have changed since the last backup. Thus, subsequent backups should take less time than the initial backup. The time required for subsequent backups depends on the amount of data accumulated, how many files have changed, and to what extent the files have changed since the last backup.

Availability of Services During Backups

Backing up a Security Management appliance copies the active data set from the 'source' Security Management appliance to a 'target' Security Management appliance with minimum disruption on the originating 'source' appliance.

The phases of the backup process and their effect on the availability of services are as follows:

- Phase 1—Phase 1 of the backup process starts with the data transfer between the source and target appliances. During data transfer, services on the source appliance remain running, therefore data collection can still continue. However, services are shut down on the target appliance. Once the data transfer is complete from the source to target appliance, Phase 2 begins.
- Phase 2—When Phase 2 begins, services on the source appliance are shut down. Any differences that have collected during the data transfer between the source and target appliance since the initial shutdown are copied to the target appliance and services on both the source and the target appliances are returned to the state they were in when backup was initiated. This allows maintain maximum uptime on the source appliance and no data loss for either appliance.

During the backup, data availability reports may not work, and when viewing the message tracking results, the hostname for each message may be labeled as 'unresolved'.

If you try to schedule a report and forget that a backup is in progress, you can check the system status by choosing **Management Appliance > Centralized Services**. From this window you can see the warning at the top of the page that a system backup is in progress.

Interruption of a Backup Process

Note If there is an unexpected reboot of the source appliance while a backup is being performed, the target appliance is unaware of this stoppage. You must cancel the backup on the target appliance.

If there is an interruption of the backup process and the backup process is not completed, the next time a backup is attempted, the Security Management appliance can start the backup process up from where it was stopped.

Canceling a backup in progress is not recommended, as the existing data will be incomplete and may not be usable until a subsequent backup is completed, especially if you receive an error. If you must cancel a backup in progress, be sure to run a complete backup as soon as possible to ensure that you always have a usable current backup.

Prevent the Target Appliance From Pulling Data Directly from Managed Appliances

Step 1 Access the command-line interface of the target appliance. For instructions, see Accessing the Command Line Interface, on page 15.

Step 2 Run the `suspendtransfers` command.

Step 3 Wait for the prompt to reappear.

Step 4 Run the `suspend` command.

Step 5 Wait for the prompt to reappear.

Step 6 Exit the command-line interface of the target appliance.

Receiving Alerts About Backup Status

To receive alerts when backups are complete and be informed of any issues, configure the appliance to send you alerts of type System, severity Info. See Managing Alerts, on page 312.

Scheduling Single or Recurring Backups

You can schedule a single or recurring backup to occur at a predetermined time.

Note A backup process will not start if there are any ongoing backups on the remote machine.

Before You Begin

- Address the items in Restrictions and Requirements for Backups , on page 290.
- Make sure you disable two-factor authentication temporarily on the target appliance before you begin the backup process. Once the backup process is complete, you can enable two-factor authentication on the target appliance.

Step 1 Login, as administrator, to the command-line interface of the source appliance.

Step 2 At the command prompt, type **backupconfig**and press **Enter**.

Step 3 If the connection between source and target appliances is slow, turn on data compression:
Type **setup** and enter **Y**.

Step 4 Type **Schedule** and press **Enter**.

Step 5 Type the IP address of the target Security Management appliance.

Step 6 Enter a meaningful name to identify the target appliance (up to 20 characters).

Step 7 Enter the admin user name and password for the target appliance.

Step 8 Respond to prompts about which data you want to back up.

Step 9 To schedule a single backup, type **2** to Schedule a single backup and press **Enter**.

Step 10 To schedule a recurring backup:

a) Type **1** to Setup Repeating Backup Schedule and press **Enter**.

b) Choose the frequency for your periodic backup and press **Enter**.

Step 11 Type the specific date or day and time that you want the backup to start and press **Enter**.

Step 12 Type the name of the backup process.

Step 13 Verify that the backup was successfully scheduled: Type **View** and press **Enter** at the command prompt.

Step 14 See also Other Important Backup Tasks , on page 295.

Starting an Immediate Backup

Note A backup process will not start if there are any ongoing backups on the target machine.

Before You Begin

Meet all requirements in Restrictions and Requirements for Backups , on page 290.

Step 1 Login, as administrator, to the command-line interface of the source appliance.

Step 2 At the command prompt, type **backupconfig** and press **Enter**.

Step 3 If the connection between source and target appliances is slow, turn on data compression:
Type **setup** and enter **Y**.

Step 4 Type **Schedule** and press **Enter**.

Step 5 Type the IP address of the target Security Management appliance.

Step 6 Enter a meaningful name to identify the target appliance (up to 20 characters).

Step 7 Enter the admin user name and password for the target appliance.

Step 8 Respond to prompts about which data you want to back up.

Step 9 Type **3**to Start a Single Backup Now and press **Enter**.

Step 10 Enter a meaningful name for the backup job.
The backup process begins in a few minutes.

Step 11 (Optional) To see the progress of the backup, type **Status** at the command-line prompt.

Step 12 See also Other Important Backup Tasks , on page 295.

Checking Backup Status

Step 1 Log in, as administrator, to the command-line interface of the primary appliance.

Step 2 At the command prompt, type **backupconfig** and press **Enter**.

To Check Status Of	Do This
A scheduled backup	Choose the View operation.
A backup in progress	Choose the Status operation. If you have configured alerts, check your email or see Viewing Recent Alerts , on page 314.

Backup Information in Log Files

Backup logs record the backup process from start to finish.

Information about backup scheduling is in the SMA logs.

Related Topics

- Checking Backup Status , on page 295

Other Important Backup Tasks

Consider doing the following in order to prevent loss of items that are not backed up by the backup processes described in this section, and to speed setup of your replacement Security Management appliance in case of appliance failure:

- To save the settings from your primary Security Management appliance, see Saving and Importing Configuration Settings , on page 327. Save the configuration file to a safe location separate from your primary Security Management appliance.
- To save log files from your Security Management appliance to an alternate location, see Log Subscriptions, on page 371.

Additionally, you can set up a log subscription for Backup Logs. See Creating a Log Subscription in the GUI, on page 373.

Making a Backup Appliance the Primary Appliance

If you are upgrading appliance hardware, or if you need to switch appliances for any other reason, use this procedure.

Before You Begin

Review the information in Backing Up Security Management Appliance Data , on page 290.

Step 1 Save a copy of the configuration file from your old/primary/source appliance to a location that you can reach from the new appliance. See Saving and Importing Configuration Settings , on page 327.

Step 2 Run the System Setup Wizard on the new/backup/target appliance.

Step 3 Meet the requirements in Restrictions and Requirements for Backups , on page 290.

Step 4 Run a backup from the old/primary/source appliance. See instructions at Starting an Immediate Backup, on page 294.

Step 5 Wait for the backup to complete.

Step 6 Run the suspendtransfers and suspend commands on the old/primary/source appliance.

Step 7 Run a second backup to transfer last-minute data from the old/primary/source to the new/backup/target appliance.

Step 8 Import the configuration file into the new/backup/target appliance.

Step 9 Run the resumetransfers and resume commands on the new/backup/target appliance.
Do NOT run this command on the old/original primary/source appliance.

Step 10 Establish the connection between the new/backup/target appliance and the managed email security appliances:

Step 11
a) [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.
b) Select **Management Appliance > Centralized Services > Security Appliances**.
c) Click an appliance name.
d) Click the **Establish Connection** button.
e) Click **Test Connection**.
f) Return to the list of appliances.
g) Repeat for each managed appliance.

Step 12 Verify that the new/target appliance is now functioning as the primary appliance:
Select **Management Appliance > Centralized Services > System Status** and check the status of data transfers.

Disaster Recovery on the Security Management Appliance

If your Security Management appliance unexpectedly fails, use the following procedure to restore security management services and your backed-up data, which you regularly save using the information in Backing Up Security Management Appliance Data , on page 290.

A typical appliance configuration might look as shown in the following figure:

Figure 13: Disaster Recovery: A Typical Environment

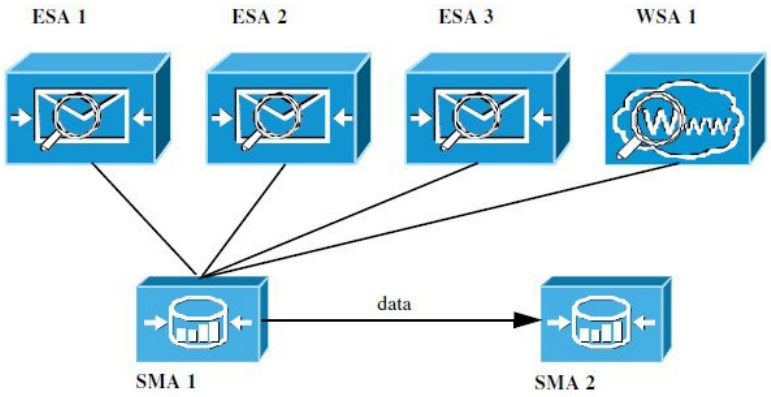

In this environment, SMA 1 is the primary Security Management appliance that is receiving data from ESAs 1-3 and WSA 1. SMA 2 is the backup Security Management appliance receiving backup data from SMA1.

In case of failure, you must configure SMA 2 to be your primary Security Management appliance.

To configure SMA 2 as your new primary Security Management appliance and restore service:

DETAILED STEPS

	Command or Action	Purpose
Step 1	If you are using Centralized Policy, Virus, and Outbreak Quarantines: • On each Email Security appliance, disable the centralized quarantines.	See instructions for disabling Centralized Policy, Virus, and Outbreak Quarantines in the Email Security appliance documentation. This will create local quarantines on each Email Security appliance, which you will migrate later to the new Security Management appliance.
Step 2	Load onto your backup Security Management appliance (SMA2) the configuration file that you saved from your primary Security Management appliance (SMA1).	See Loading a Configuration File, on page 328.
Step 3	Recreate the IP address from the failed SMA 1 to be the IP address on SMA 2	**1** On SMA 2 choose **Network > IP Interfaces > Add IP Interfaces**. **2** On the **Add IP Interfaces** page, enter all of the relevant IP Interface information from the failed SMA1 into the text fields to recreate the interface on SMA 2. For more information about Adding IP Interfaces, see Configuring IP Interfaces, on page 390.
Step 4	Submit and commit your changes.	

	Command or Action	Purpose
Step 5	Enable all applicable centralized services on the new Security Management appliance (SMA 2).	See Configuring Services on the Security Management Appliance, on page 20.
Step 6	Add all appliances on to the new Security Management appliance (SMA 2). • Test to see that each appliance is enabled and working by establishing a connection to the appliances and testing the connections.	See About Adding Managed Appliances, on page 19.
Step 7	If you are using Centralized Policy, Virus, and Outbreak Quarantines, configure quarantine migration on the new Security Management appliance, then enable and configure the migration on each applicable Email Security appliance.	See Centralizing Policy, Virus, and Outbreak Quarantines , on page 199.
Step 8	If necessary, restore additional data.	See Other Important Backup Tasks , on page 295.

What to Do Next

After this process is complete, SMA 2 becomes the primary Security Management appliance. All data from ESAs 1-3 and WSA 1 now goes to SMA 2, as shown in the following figure:

Figure 14: Disaster Recovery: Final Result

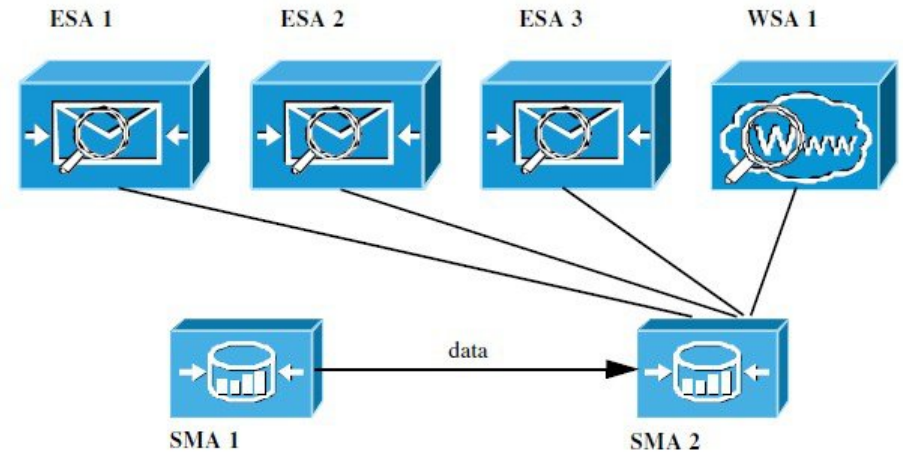

Upgrading Appliance Hardware

See Making a Backup Appliance the Primary Appliance , on page 295.

Upgrading AsyncOS

Batch Commands for Upgrades

Batch commands for upgrade procedures are documented in the CLI Reference Guide for AsyncOS for Email at
http://www.cisco.com/c/en/us/support/security/email-security-appliance/products-command-reference-list.html

Determining Network Requirements for Upgrades and Updates

The update servers for Cisco content security appliances use dynamic IP addresses. If you have strict firewall policies, you may need to configure a static location for AsyncOS upgrades. If you determine that your firewall settings require a static IP for upgrades, contact Cisco Customer support to obtain the required URL addresses.

Note If you have any existing firewall rules allowing download of legacy upgrades from upgrades.cisco.com ports such as 22, 25, 80, 4766, they will need to be removed and/or replaced with revised firewall rules.

Choosing an Upgrade Method: Remote vs. Streaming

Cisco provides two methods (or 'sources') for upgrading AsyncOS on your appliances:

- Streaming upgrades — Each appliance downloads the AsyncOS upgrades via HTTP directly from the Cisco content security update servers.
- Remote upgrades — You only download the upgrade image from Cisco one time, and then serve it to your appliances. Your appliances then download the AsyncOS upgrades from a server within your network.

You will configure the upgrade method in Configuring Upgrade and Service Update Settings, on page 302. Optionally, use the **updateconfig** command in the CLI.

Streaming Upgrade Overview

In Streaming upgrades, each Cisco Content Security appliance connects directly to the Cisco content security update servers to find and download upgrades:

Figure 15: Streaming Update Method

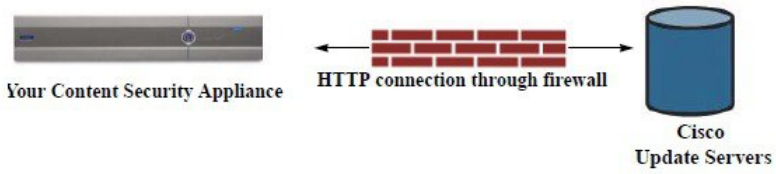

This method requires that your appliance contacts the Cisco content security update servers directly from the network.

Remote Upgrade Overview

You can also download and host updates to AsyncOS locally from within your own network (Remote Upgrade) rather than obtaining updates directly from the Cisco update servers (Streaming Upgrades). Using this feature, an encrypted update image downloaded via HTTP to any server in your network that has access to the Internet. If you choose to download the update image, you can then configure an internal HTTP server (an "update manager") to host the AsyncOS images to your Security Management appliances.

Figure 16: Remote Update Method

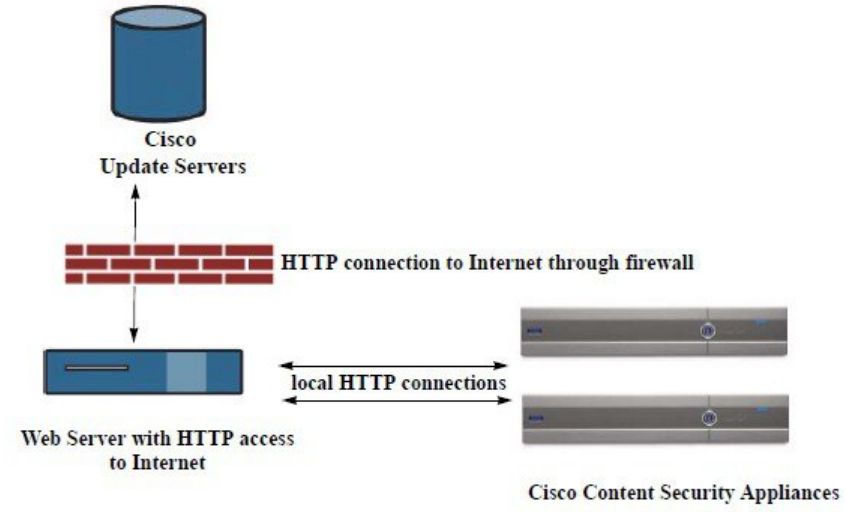

The basic process is as follows:

Step 1 Read the information in Hardware and Software Requirements for Remote Upgrades, on page 301 and Hosting a Remote Upgrade Image, on page 301.

Step 2 Configure a local server to retrieve and serve the upgrade files.

Step 3 Download the upgrade files.

Step 4 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 5 Choose **Management Appliance > System Administration > Update Settings**
From this page, choose to configure the appliance to use the local server.

Step 6 Choose **Management Appliance > System Administration > System Upgrade**

Step 7 Click **Available Upgrades**.

Note From the command-line prompt you can also do the following: Run the **updateconfig** command then run the **upgrade** command.

For complete information, see Upgrading AsyncOS, on page 299.

Hardware and Software Requirements for Remote Upgrades

For *downloading* AsyncOS upgrade files, you must have a system in your internal network that has:

- Internet access to the update servers for Cisco content security appliances.
- A web browser.

Note For this release, if you need to configure a firewall setting to allow HTTP access to this address, you must configure it using the DNS name and not a specific IP address.

For *hosting* AsyncOS update files, you must have a server in your internal network that has:

- A web server — for example, Microsoft IIS (Internet Information Services) or the Apache open source server — that:
  - supports the display of directory or filenames in excess of 24 characters
  - has directory browsing enabled
  - is configured for anonymous (no authentication) or basic ("simple") authentication
  - contains at least 350MB of free disk space for each AsyncOS update image

Hosting a Remote Upgrade Image

After setting up a local server, go to **http://updates.ironport.com/fetch_manifest.html** to download a zip file of an upgrade image. To download the image, enter your serial number and the version number of the Cisco Content Security appliance. You will then be presented with a list of available upgrades. Click the upgrade version that you want to download a zip file of the upgrade image. To use the upgrade image for AsyncOS upgrades, enter the base URL for your local server on the Edit Update Settings page (or use updateconfig in the CLI).

You can also host an XML file on a local server that limits the available upgrades for the Cisco Content Security appliances on your network to the version selected at http://updates.ironport.com/fetch_manifest.html . Your Cisco Content Security appliances still download the upgrade from the Cisco servers. If you want to host the upgrade list on a local server, download the zip file and extract the asyncos/phoebe-my-upgrade.xml file to the root directory of the local server. To use the upgrade list for AsyncOS upgrades, enter the full URL for the XML file on the Edit Update Settings page (or use updateconfig in the CLI).

For more information about remote upgrades, check the Knowledge Base (see Knowledge Base Articles (TechNotes) , on page 407) or contact your support provider.

Important Differences in Remote Upgrading Method

Note these differences when upgrading AsyncOS from a local server (Remote upgrade) as opposed to the Streaming upgrade method:

- The upgrade installs immediately *while downloading* .

- A banner appears for 10 seconds at the beginning of the upgrade process. While this banner appears, you have the option to press Control-C to exit the upgrade process before downloading starts.

Configuring Upgrade and Service Update Settings

You can configure how the Cisco Content Security appliance downloads security services updates (such as time zone rules) and AsyncOS upgrades. For example, you can choose whether to download upgrades and updates dynamically from Cisco servers or from a local server onto which you have made the images available; configure the update interval; or disable automatic updates.

AsyncOS periodically queries the update servers for new updates to all security service components except for new AsyncOS upgrades. To upgrade AsyncOS, you must manually prompt AsyncOS to query for available upgrades.

You can configure upgrade and updates settings in the GUI (see the following two sections) or using the updateconfig command in the CLI.

You can also configure upgrade notification settings.

Upgrade and Update Settings

The following table describes the update and upgrade settings you can configure.

Table 53: Update Settings for Security Services

Setting	Description
Update Servers (images)	Choose whether to download AsyncOS upgrade and service update software images, such as time zone rules and Feature Key updates, from the Cisco servers or a from a local web server. The default is the Cisco servers for both upgrades and updates. You might want to use a local web server if : • You need to download images to your appliance from a static address. See Static Upgrade and Update Server Settings for Environments with Strict Firewall Policies , on page 304. • You want to download AsyncOS upgrade images to your appliance at your convenience. (You can still download service update images dynamically from the Cisco update servers.) When you choose a local update server, enter the base URL and port number for the servers used to download the upgrades and updates. If the server requires authentication, you can also enter a valid user name and password. For more information, see Choosing an Upgrade Method: Remote vs. Streaming , on page 299 and Remote Upgrade Overview, on page 300.

Setting	Description
Update Servers (lists)	Choose whether to download the lists of available upgrades and service updates (the manifest XML files) from the Cisco servers or from a local web server. The default for both upgrades and updates is the Cisco servers. You can choose different settings for upgrades and for updates. If applicable, see Static Upgrade and Update Server Settings for Environments with Strict Firewall Policies , on page 304. If you choose local update servers, enter the full path to the manifest XML file for each list including the file name and port number for the server. If you leave the port field blank, AsyncOS uses port 80. If the server requires authentication, you can also enter a valid user name and password. For more information, see Choosing an Upgrade Method: Remote vs. Streaming , on page 299 and Remote Upgrade Overview, on page 300.
Automatic Updates	Choose whether or not to enable automatic updates for time zone rules. When enabled, enter the time to wait between checks for updates. Add a trailing **m** for minutes, **h** for hours, and **d** for days.
Interface	Choose which network interface to use when contacting the update servers for time zone rules and AsyncOS upgrades. The available proxy data interfaces are shown. By default, the appliance selects an interface to use.
HTTP Proxy Server	If an upstream HTTP proxy server exists and requires authentication, enter the server information and user name and password here. Note that if you specify a proxy server, it will be used to access and update the services listed in the GUI. This proxy server is also used to obtain File Analysis report details from the cloud
HTTPS Proxy Server	If an upstream HTTPS proxy server exists and requires authentication, enter the server information and user name and password here. Note that if you specify a proxy server, it will be used to access and update the services listed in the GUI. This proxy server is also used to obtain File Analysis report details from the cloud.

Static Upgrade and Update Server Settings for Environments with Strict Firewall Policies

The AsyncOS update servers use dynamic IP addresses. If your environment has strict firewall policies which require static IP addresses, use the following settings on the Update Settings page:

Figure 17: Static URLs for Update Servers (images) Settings

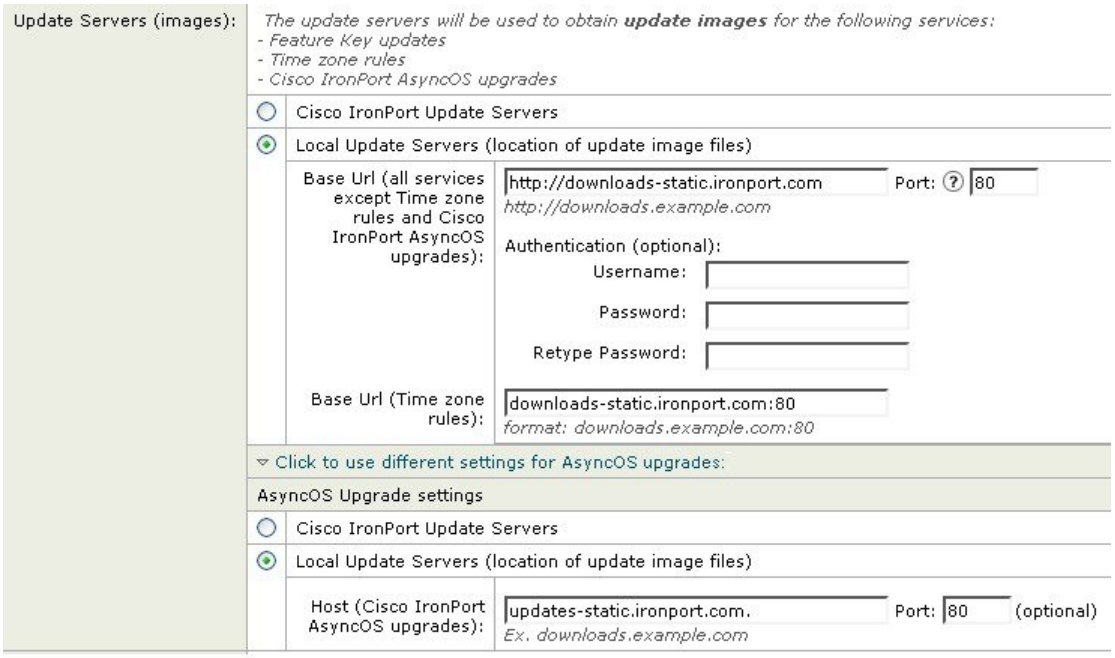

Figure 18: Static URLs for Update Servers (list) Settings

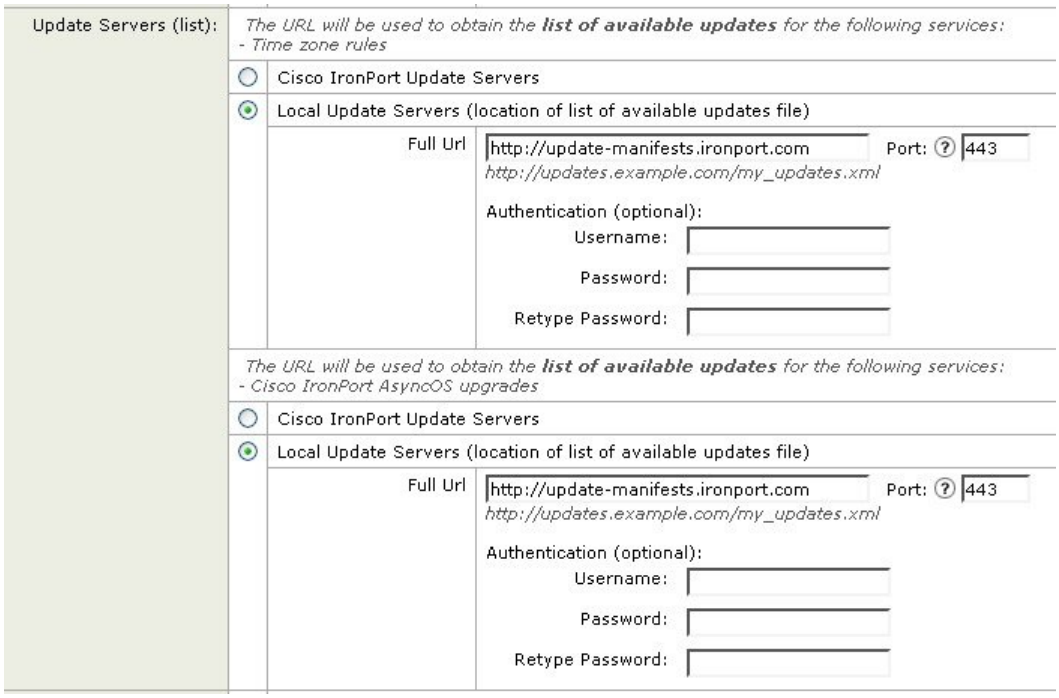

Table 54: Static Addresses for Environments with Strict Firewall Policies

Section	Setting	Static URL/IP Address and Port
Update Servers (images):	Base URL (all services except Time zone rules and AsyncOS upgrades)	http://downloads-static.ironport.com 204.15.82.8 Port 80
	Base URL (Time zone rules)	downloads-static.ironport.com 204.15.82.8 Port 80
	Host (AsyncOS upgrades)	updates-static.ironport.com 208.90.58.25 Port 80

Section	Setting	Static URL/IP Address and Port
Update Servers (list):	For updates on physical hardware appliances: Full URL	update-manifests.ironport.com 208.90.58.5 Port 443
	For updates on virtual appliances: Full URL	update-manifests.sco.cisco.com Port 443
	For upgrades: Full URL	update-manifests.ironport.com 208.90.58.5 Port 443

Configuring the Update and Upgrade Settings from the GUI

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > Update Settings.**

Step 3 Click **Edit Update Settings**.
Use the descriptions in Upgrade and Update Settings, on page 302 to configure the settings in this procedure.

Step 4 In the **Update Servers (images)** section, specify the servers from which to download images for updates.

Step 5 Specify the server from which to download images for AsyncOS upgrades:

a) At the bottom of the same section, click the **Click to use different settings for AsyncOS upgrades**link.

b) Specify server settings for downloading images for AsyncOS upgrades.

Step 6 In the **Update Servers (list)** section, specify the servers for obtaining the list of available updates and AsyncOS upgrades. The top subsection applies to updates. The bottom subsection applies to upgrades.

Step 7 Specify settings for Time Zone rules and interface.

Step 8 (Optional) Specify settings for Proxy Servers.

Step 9 Submit and commit your changes.

Step 10 Verify that your results are what you expect:
If you are not already looking at the Update Settings page, choose **Management Appliance > System Administration > Update Settings.**

Some URLs may append an “asyncos” directory to the server URL. You can ignore this discrepancy.

Upgrade Notifications

By default, users with administrator and technician privileges will see a notification at the top of the web interface when an AsyncOS upgrade is available for the appliance.

To	Do This
View more information about the latest upgrade	Hover over the upgrade notification.
View a list of all available upgrades	Click the down arrow in the notification.
Dismiss a current notification. The appliance will not display another notification until a new upgrade becomes available.	Click the down arrow, then select **Clear the notification**, then click **Close**.
Prevent future notifications (Users with Administrator privileges only.)	Go to **Management Appliance > System Administration > System Upgrade.**

Before You Upgrade: Important Steps

Before You Begin

See network requirements at Determining Network Requirements for Upgrades and Updates , on page 299.

Step 1 Take steps to prevent or minimize data loss:

- Make sure the new appliance has sufficient disk capacity and the same or greater size allocations for each data type that will be transferred. See About Disk Space Maximums and Allocations, on page 334.
- If you have received any disk space warnings, resolve any disk space issues before upgrading.

Step 2 Save the XML configuration file off the appliance. See caveats at Saving and Exporting the Current Configuration File, on page 328.
If you need to revert to the pre-upgrade release for any reason, you will need this file.

Step 3 If you are using the Safelist/Blocklist feature, export the list off the appliance.
Click **Management Appliance > System Administration > Configuration File** and scroll down.

Step 4 Suspend the listeners using the **suspendlistener** command when running the upgrade from the CLI. If you perform the upgrade from the GUI, listener suspension occurs automatically.

Step 5 Drain the mail queue and the delivery queue.

Step 6 Verify that the upgrade settings are configured as you want them. See Configuring Upgrade and Service Update Settings, on page 302.

Upgrading AsyncOS

You can download and install in a single operation, or download in the background and install later.

Note When downloading and upgrading AsyncOS in a single operation from a local server instead of from a Cisco server, the upgrade installs immediately *while downloading* . A banner displays for 10 seconds at the beginning of the upgrade process. While this banner is displayed, you have the option to type Control-C to exit the upgrade process before downloading starts.

Before You Begin

- Choose whether you will download upgrades directly from Cisco or will host upgrade images from a server on your network. Then set up your network to support the method you choose. Then configure the appliance to obtain upgrades from your chosen source. See Choosing an Upgrade Method: Remote vs. Streaming , on page 299 and Configuring Upgrade and Service Update Settings, on page 302.
- Before installing the upgrade, follow the instructions in Before You Upgrade: Important Steps , on page 307.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > System Upgrade**.

Step 3 Click **Upgrade Options**.

Step 4 Choose an option:

To	Do This
Download and install the upgrade in a single operation	Click **Download and Install**. If you have already downloaded an installer, you will be prompted to overwrite the existing download.
Download an upgrade installer	Click **Download only**. If you have already downloaded an installer, you will be prompted to overwrite the existing download. The installer downloads in the background without interrupting service.
Install a downloaded upgrade installer	Click **Install**. This option appears only if an installer has been downloaded. The AsyncOS version to be installed is noted below the Install option.

Step 5 Unless you are installing a previously-downloaded installer, select an AsyncOS version from the list of available upgrades.

Step 6 If you are installing:

a) Choose whether or not to save the current configuration to the configuration directory on the appliance.

b) Choose whether or not to mask the passwords in the configuration file.

Note You cannot load a configuration file with masked passwords using the Configuration File page in the GUI or the loadconfig command in the CLI.

c) If you want to email copies of the configuration file, enter the email addresses to which you want to email the file. Use commas to separate multiple email addresses.

Step 7 Click **Proceed**.

Step 8 If you are installing:

a) Be prepared to respond to prompts during the process.
The process pauses until you respond.

A progress bar appears near the top of the page.

b) At the prompt, click **Reboot Now**.

Note Do not interrupt power to the appliance for any reason (even to troubleshoot an upgrade issue) until at least 20 minutes have passed since you rebooted.

c) After about 10 minutes, access the appliance again and log in.

What to Do Next

- If the process was interrupted, you must start the process again.
- If you downloaded but did not install the upgrade:

  When you are ready to install the upgrade, follow these instructions from the beginning, including the prerequisites in the Before You Begin section, but choose the Install option.
- If you installed the upgrade, see After Upgrading , on page 310.

Viewing Status of, Canceling, or Deleting a Background Download

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > System Upgrade**.

Step 3 Click **Upgrade Options**.

Step 4 Choose an option:

To	Do This
View download status	Look in the middle of the page. If there is no download in progress and no completed download waiting to be installed, you will not see download status information. Upgrade status also appears in upgrade_logs.

To	Do This
Cancel a download	Click the **Cancel Download** button in the middle of the page. This option appears only while a download is in progress.
Delete a downloaded installer	Click the **Delete File** button in the middle of the page. This option appears only if an installer has been downloaded.

After Upgrading

After the upgrade is complete, complete the following:

- (For deployments with associated Email Security appliances) Re-enable the listeners.
- Consider saving your configuration. For more information, see Saving and Importing Configuration Settings , on page 327.
- Before viewing the online help after upgrade, clear your browser cache, exit the browser, then open it again. This clears the browser cache of any outdated content.

About Reverting to an Earlier Version of AsyncOS

You can revert to an to a previous qualified version of AsyncOS for emergency uses.

You can also revert to the currently running build if you want to clear all data on the appliance and start with a new, clean configuration.

Related Topics

Important Note About Reversion Impact

Using the revert command on a Cisco Content Security appliance is a very destructive action. This command permanently destroys all existing configurations and data. In addition, it disrupts mail handling until the appliance is reconfigured.

Reverting does not affect feature key or virtual appliance license expiration dates.

Reverting AsyncOS

Before You Begin

- Back up or save any data that you want to preserve to a location off the appliance.

- You must have a configuration file for the version you want to revert to. Configuration files are *not* backwards-compatible.
- Because this command destroys all configuration, it is highly recommended that you have physical local access to the appliance when reverting.
- If quarantines are enabled on your Email Security appliances, disable centralization so that messages are quarantined locally on those appliances.

Step 1 Ensure that you have the configuration file for the version you want to revert to. Configuration files are not backwards-compatible.

Step 2 Save a backup copy of the current configuration of your appliance (with passwords unmasked) on another machine. To do this, you can email the file to yourself or FTP the file. A simple way to do this is to run the `mailconfig` CLI command, which emails the current configuration file on your appliance to the specified email address.

Note This is not the configuration file you will load after reverting.

Step 3 If you use the Safelist/Blocklist feature, export the Safelist/Blocklist database to another machine.

Step 4 Suspend any listeners on your Email Security appliances.

Step 5 Wait for the mail queue to empty.

Step 6 Log in to the CLI of the appliance you want to revert.
When you run the revert command, several warning prompts are issued. Once these warning prompts are accepted, the revert action takes place immediately. Therefore, do not begin the reversion process until after you have completed the prereversion steps.

Step 7 From the command-line prompt, type the **`revert`** command and respond to the prompts.
The following example shows the **`revert`** command:

Example:

```
m650p03.prep> revert
This command will revert the appliance to a previous version of AsyncOS.
WARNING: Reverting the appliance is extremely destructive.
The following data will be destroyed in the process:
- all configuration settings (including listeners)
- all log files
- all databases (including messages in Virus Outbreak and Policy
quarantines)
- all reporting data (including saved scheduled reports)
- all message tracking data
- all Cisco Spam Quarantine message and end-user safelist/blocklist data
Only the network settings will be preseved.
Before running this command, be sure you have:
- saved the configuration file of this appliance (with passwords
unmasked)
- exported the Cisco Spam Quarantine safelist/blocklist database
   to another machine (if applicable)
- waited for the mail queue to empty
Reverting the device causes an immediate reboot to take place.
After rebooting, the appliance reinitializes itself and reboots again to the desired version.
Do you want to continue? yes
Are you sure you want to continue? yes
Available versions
=================
 1. 7.2.0-390
 2. 6.7.6-020
Please select an AsyncOS version: 1
You have selected "7.2.0-390".
```

```
Reverting to "testing" preconfigure install mode.
The system will now reboot to perform the revert operation.
```

Step 8 Wait for the appliance to reboot twice.

Step 9 Log in to the appliance using the CLI.

Step 10 Add at least one appliance and wait a few minutes to allow any URL Category updates to be downloaded from that appliance.

Step 11 After URL Category updates are completed, load the XML configuration file of the version you are reverting to.

Step 12 If you use the Safelist/Blocklist feature, import and restore the Safelist/Blocklist database.

Step 13 Reenable any listeners on your Email Security appliances.

Step 14 Commit your changes.
The reverted Cisco Content Security appliance should now run using the selected AsyncOS version.

Note It may take 15 to 20 minutes before reversion is complete and console access to the Cisco Content Security appliance is available again.

About Updates

Service updates are periodically made available for download. To specify settings for these downloads, see Configuring Upgrade and Service Update Settings, on page 302

Related Topics

Configuring the Return Address for Generated Messages

You can configure the envelope sender for mail generated by AsyncOS for the following types of cases:

- Bounce messages
- Reports

You can specify the display, user, and domain names of the return address. You can also choose to use the Virtual Gateway domain for the domain name.

Use the Return Addresses page available on the System Administration menu in the GUI, or use the **addressconfig** command in the CLI.

To modify the return address for system-generated email messages in the GUI, click **Edit Settings** on the Return Addresses page. Make changes to the address or addresses you want to modify, click **Submit,** and commit your changes.

Managing Alerts

The appliance sends you email alerts about events occurring on the appliance.

To	Do This
Have different types of alerts sent to different administrative users	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface. Select **Management Appliance > System Administration > Alerts** If you enabled AutoSupport during system setup, the email address that you specified will receive alerts for all severities and classes by default. You can change the configuration at any time. Separate multiple addresses with commas.
Configure global settings for alerts, including: • Alert sender (FROM:) address • Controls for duplicate alerts • AutoSupport settings.	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface. Select **Management Appliance > System Administration > Alerts** See About Duplicate Alerts , on page 315 See Cisco AutoSupport, on page 315
View a list of recent alerts Manage settings for this list	See Viewing Recent Alerts , on page 314
See a list of alerts and their descriptions	See: Hardware Alert Descriptions , on page 315. System Alert Descriptions , on page 316
Understand alert delivery mechanisms	See Alert Delivery, on page 314

Alert Types and Severities

Alert types include:

- Hardware alerts. See Hardware Alert Descriptions , on page 315.
- System alerts. See System Alert Descriptions , on page 316.
- Updater alerts.

Alerts can have the following severities:

- Critical: issue that requires immediate attention
- Warning: problem or error requiring further monitoring and potentially immediate attention

- Info: information generated in the routine functioning of this device

Alert Delivery

Because alert messages can be used to inform you of problems within your Cisco Content Security appliance, they are not sent using AsyncOS's normal mail delivery system. Instead, alert messages pass through a separate and parallel email system designed to operate even in the face of significant system failure in AsyncOS.

The alert mail system does not share the same configuration as AsyncOS, which means that alert messages may behave slightly differently from other mail delivery:

- Alert messages are delivered using standard DNS MX and A record lookups.
  - They do cache the DNS entries for 30 minutes and the cache is refreshed every 30 minutes, so in case of DNS failure the alerts still go out.
- If your deployment includes Email Security appliances:
  - Alert messages do not pass through the work queue, so they are not scanned for viruses or spam. They are also not subjected to message filters or content filters.
  - Alert messages do not pass through the delivery queue, so they will not be affected by bounce profiles or destination control limits.

Viewing Recent Alerts

To	Do This
View a list of recent alerts	Users with administrator and operator access can choose **Management Appliance > System Administration > Alerts** and click the **View Top Alerts** button. Alerts appear even if there was a problem emailing them.
Sort the list	Click a column heading.
Specify the maximum number of alerts to save in this list	Use the alertconfig command in the command-line interface
Disable this feature	Use the alertconfig command in the command-line interface to set the maximum number of alerts to zero (0).

About Duplicate Alerts

You can specify the initial number of seconds to wait before AsyncOS will send a duplicate alert. If you set this value to 0, duplicate alert summaries are not sent; instead, all duplicate alerts are sent without any delay (this can lead to a large amount of email over a short amount of time). The number of seconds to wait between sending duplicate alerts (alert interval) is increased after each alert is sent. The increase is the number of seconds to wait plus twice the last interval. So a 5-second wait would have alerts sent at 5 seconds, 15 seconds, 35 seconds, 75 seconds, 155 seconds, 315 seconds, and so on.

Eventually, the interval could become large. You can set a cap on the number of seconds to wait between intervals via the maximum number of seconds to wait before sending a duplicate alert field. For example, if you set the initial value to 5 seconds, and the maximum value to 60 seconds, alerts would be sent at 5 seconds, 15 seconds, 35 seconds, 60 seconds, 120 seconds, and so on.

Cisco AutoSupport

To allow Cisco to better support and design future system changes, the Cisco Content Security appliance can be configured to send Cisco a copy of all alert messages generated by the system. This feature, called 'AutoSupport', is a useful way to allow Customer Support to be proactive in supporting your needs. AutoSupport also sends weekly reports noting the uptime of the system, the output of the **status** command, and the AsyncOS version used.

By default, alert recipients set to receive Information severity level alerts for System alert types receive a copy of every message sent to Cisco. This can be disabled if you do not want to send the weekly alert messages internally. To enable or disable this feature, select **Management Appliance > System Administration Alerts** and click edit settings.

By default, if AutoSupport is enabled, the weekly AutoSupport report is sent to alert recipients set to receive system alerts at the Information level.

Hardware Alert Descriptions

Table 55: Hardware Alert Descriptions

Alert Name	Description	Severity
INTERFACE.ERRORS	Sent when interface errors are detected.	Warning
MAIL.MEASUREMENTS_FILESYSTEM	Sent when a disk partition is nearing capacity (75%).	Warning
MAIL.MEASUREMENTS_FILESYSTEM.CRITICAL	Sent when a disk partition reaches 90% capacity (and at 95%, 96%, 97%, and so on).	Critical
SYSTEM.RAID_EVENT_ALERT	Sent when a critical RAID-event occurs.	Warning
SYSTEM.RAID_EVENT_ALERT_INFO	Sent when a RAID-event occurs.	Information

System Alert Descriptions

Table 56: System Alert Descriptions

Alert Name	Description	Severity
COMMON.APP_FAILURE	Sent when there is an unknown application failure.	Critical
COMMON.KEY_EXPIRED_ALERT	Sent when a feature key has expired.	Warning
COMMON.KEY_EXPIRING_ALERT	Sent when a feature key is about to expire.	Warning
COMMON.KEY_FINAL_EXPIRING_ALERT	Sent as a final notice that a feature key is about to expire.	Warning
DNS.BOOTSTRAP_FAILED	Sent when the appliance is unable to contact the root DNS servers.	Warning
COMMON.INVALID_FILTER	Sent when an invalid filter is encountered.	Warning

Alert Name	Description	Severity
IPBLOCKD.HOST_ADDED_TO_WHITELIST IPBLOCKD.HOST_ADDED_TO_BLACKLIST IPBLOCKD.HOST_REMOVED_FROM_BLACKLIST	Alert messages: • The host at <IP address> has been added to the blacklist because of an SSH DOS attack. • The host at <IP address> has been permanently added to the ssh whitelist. • The host at <IP address> has been removed from the blacklist IP addresses that try to connect to the appliance over SSH but do not provide valid credentials are added to the SSH blacklist if more than 10 failed attempts occur within two minutes. When a user logs in successfully from the same IP address, that IP address is added to the whitelist. Addresses on the whitelist are allowed access even if they are also on the blacklist.	Warning
LDAP.GROUP_QUERY_FAILED_ALERT	Sent when an LDAP group query fails.	Critical
LDAP.HARD_ERROR	Sent when an LDAP query fails completely (after trying all servers).	Critical
LOG.ERROR.*	Various logging errors.	Critical
MAIL.PERRCPT.LDAP_GROUP_QUERY_FAILED	Sent when an LDAP group query fails during per-recipient scanning.	Critical
MAIL.QUEUE.ERROR.*	Various mail queue hard errors.	Critical
MAIL.RES_CON_START_ALERT.MEMORY	Sent when RAM utilization has exceeded the system resource conservation threshold.	Critical

Alert Name	Description	Severity
MAIL.RES_CON_START_ALERT.QUEUE_SLOW	Sent when the mail queue is overloaded and system resource conservation is enabled.	Critical
MAIL.RES_CON_START_ALERT.QUEUE	Sent when queue utilization has exceeded the system resource conservation threshold.	Critical
MAIL.RES_CON_START_ALERT.WORKQ	Sent when listeners are suspended because the work queue size is too big.	Critical
MAIL.RES_CON_START_ALERT	Sent when the appliance enters “resource conservation” mode.	Critical
MAIL.RES_CON_STOP_ALERT	Sent when the appliance leaves “resource conservation” mode.	Critical
MAIL.WORK_QUEUE_PAUSED_NATURAL	Sent when the work queue is paused.	Critical
MAIL.WORK_QUEUE_UNPAUSED_NATURAL	Sent when the work queue is resumed.	Critical
NTP.NOT_ROOT	Sent when the appliance is unable to adjust time because NTP is not running as root.	Warning
PERIODIC_REPORTS.DOMAIN_REPORT. DOMAIN_FILE_ERRORS	Sent when errors are found in the domain specification file.	Critical
PERIODIC_REPORTS.DOMAIN_REPORT.FILE_EMPTY	Sent when the domain specification file is empty.	Critical
PERIODIC_REPORTS.DOMAIN_REPORT.FILE_MISSING	Sent when the domain specification file is not found.	Critical
REPORTD.DATABASE_OPEN_FAILED_ALERT	Sent if the reporting engine is unable to open the database.	Critical
REPORTD.AGGREGATION_DISABLED_ALERT	Sent if the system runs out of disk space. When the disk usage for a log entry exceeds the log usage threshold, reportd disables aggregation and sends the alert.	Warning
REPORTING.CLIENT.UPDATE_FAILED_ALERT	Sent if the reporting engine was unable to save reporting data.	Warning

Alert Name	Description	Severity
REPORTING.CLIENT.JOURNAL.FULL	Sent if the reporting engine is unable to store new data.	Critical
REPORTING.CLIENT.JOURNAL.FREE	Sent when the reporting engine is again able to store new data.	Information
PERIODIC_REPORTS.REPORT_TASK. BUILD_FAILURE_ALERT	Sent when the reporting engine is unable to build a report.	Critical
PERIODIC_REPORTS.REPORT_TASK. EMAIL_FAILURE_ALERT	Sent when a report could not be emailed.	Critical
PERIODIC_REPORTS.REPORT_TASK. ARCHIVE_FAILURE_ALERT	Sent when a report could not be archived.	Critical
SENDERBASE.ERROR	Sent when an error occurred while processing a response from SenderBase.	Information
SMAD.ICCM.ALERT_PUSH_FAILED	Sent if a configuration push failed for one or more hosts.	Warning
SMAD.TRANSFER.TRANSFERS_STALLED	Sent if SMA logs are unable to fetch tracking data for two hours or reporting data for six hours.	Warning
SMTPAUTH.FWD_SERVER_FAILED_ALERT	Sent when the SMTP Authentication forwarding server is unreachable.	Warning
SMTPAUTH.LDAP_QUERY_FAILED	Sent when an LDAP query fails.	Warning
SYSTEM.HERMES_SHUTDOWN_FAILURE. REBOOT	Sent when there was a problem shutting down the system on reboot.	Warning
SYSTEM.HERMES_SHUTDOWN_FAILURE. SHUTDOWN	Sent when there was a problem shutting down the system.	Warning
SYSTEM.RCPTVALIDATION.UPDATE_FAILED	Sent when a recipient validation update failed.	Critical
SYSTEM.SERVICE_TUNNEL.DISABLED	Sent when a tunnel created for Cisco Support Services is disabled.	Information

Alert Name	Description	Severity
SYSTEM.SERVICE_TUNNEL.ENABLED	Sent when a tunnel created for Cisco Support Services is enabled.	Information

Changing Network Settings

This section describes the features used to configure the network operation of the appliance. These features give you direct access to the hostname, DNS, and routing settings that you configured using the System Setup Wizard in Running the System Setup Wizard, on page 16.

The following features are described:

- sethostname
- DNS configuration (in the GUI and by using the dnsconfig command in the CLI)
- Routing configuration (in the GUI and by using the routeconfig and setgateway commands in the CLI)
- dnsflush
- Password

Changing the System Hostname

The hostname is used to identify the system at the CLI prompt. You must enter a fully qualified hostname. The s ethostname command sets the name of the content security appliance. The new hostname does not take effect until you issue the commit command.

The sethostname Command

```
oldname.example.com> sethostname
[oldname.example.com]> mail3.example.com
oldname.example.com>
```

For the hostname change to take effect, you must enter the commit command. After you have successfully committed the hostname change, the new name appears in the CLI prompt:

```
oldname.example.com> commit
Please enter some comments describing your changes:
[]> Changed System Hostname
Changes committed: Mon Jan 04 12:00:01 2010
```

The new hostname appears in the prompt as follows: mail3.example.com>

Configuring Domain Name System Settings

You can configure the Domain Name System (DNS) settings for your content security appliance through the Management Appliance > Network > DNS page in the GUI, or via the dnsconfig command.

You can configure the following settings:

- Whether to use the Internet's DNS servers or your own, and which server(s) to use
- Which interface to use for DNS traffic
- The number of seconds to wait before timing out a reverse DNS lookup
- Clearing the DNS cache

Specifying DNS Servers

AsyncOS can use the Internet root DNS servers, your own DNS servers, or the Internet root DNS servers and authoritative DNS servers that you specify. When using the Internet root servers, you may specify alternate servers to use for specific domains. Because an alternate DNS server applies to a single domain, it must be authoritative (provide definitive DNS records) for that domain.

AsyncOS supports "splitting" DNS servers when not using the Internet's DNS servers. If you are using your own internal server, you can also specify exception domains and associated DNS servers.

When setting up "split DNS," you should set up the in-addr.arpa (PTR) entries as well. For example, if you want to redirect ".eng" queries to the nameserver 1.2.3.4 and all the .eng entries are in the 172.16 network, then you should specify "eng,16.172.in-addr.arpa" as the domains in the split DNS configuration.

Multiple Entries and Priority

For each DNS server that you enter, you can specify a numeric priority. AsyncOS attempts to use the DNS server with the priority closest to 0. If that DNS server is not responding, AsyncOS attempts to use the server at the next priority. If you specify multiple entries for DNS servers with the same priority, the system randomizes the list of DNS servers at that priority every time it performs a query. The system then waits a short amount of time for the first query to expire or "time out" and then a slightly longer amount of time for the second, and so on. The amount of time depends on the exact total number of DNS servers and priorities that have been configured. The timeout length is the same for all IP addresses at any particular priority. The first priority gets the shortest timeout; each subsequent priority gets a longer timeout. Further, the timeout period is roughly 60 seconds. If you have one priority, the timeout for each server at that priority is 60 seconds. If you have two priorities, the timeout for each server at the first priority is 15 seconds, and each server at the second priority is 45 seconds. For three priorities, the timeouts are 5, 10, 45.

For example, suppose you configure four DNS servers, with two of them at priority 0, one at priority 1, and one at priority 2:

Table 57: Example of DNS Servers, Priorities, and Timeout Intervals

Priority	Server(s)	Timeout (Seconds)
0	1.2.3.4, 1.2.3.5	5, 5
1	1.2.3.6	10
2	1.2.3.7	45

AsyncOS randomly chooses between the two servers at priority 0. If one of the priority 0 servers is down, the other is used. If both of the priority 0 servers are down, the priority 1 server (1.2.3.6) is used, and then, finally, the priority 2 (1.2.3.7) server.

The timeout period is the same for both priority 0 servers, longer for the priority 1 server, and longer still for the priority 2 server.

Using the Internet Root Servers

The AsyncOS DNS resolver is designed to accommodate the large number of simultaneous DNS connections required for high-performance email delivery.

Note If you choose to set the default DNS server to something other than the Internet root servers, that server must be able to recursively resolve queries for domains for which it is not an authoritative server.

Reverse DNS Lookup Timeout

The Cisco Content Security appliance attempts to perform a "double DNS lookup" on all remote hosts connecting to a listener for the purposes of sending or receiving email. That is, the system acquires and verifies the validity of the remote host's IP address by performing a double DNS lookup. This consists of a reverse DNS (PTR) lookup on the IP address of the connecting host, followed by a forward DNS (A) lookup on the results of the PTR lookup. The system then checks that the results of the A lookup match the results of the PTR lookup. If the results do not match, or if an A record does not exist, the system uses only the IP address to match entries in the Host Access Table (HAT). This particular timeout period applies only to this lookup and is not related to the general DNS timeout discussed in Multiple Entries and Priority, on page 321.

The default value is 20 seconds. You can disable the reverse DNS lookup timeout globally across all listeners by entering '0' as the number of seconds. If the value is set to 0 seconds, the reverse DNS lookup is not attempted, and instead the standard timeout response is returned immediately.

DNS Alert

Occasionally, an alert may be generated with the message "Failed to bootstrap the DNS cache" when an appliance is rebooted. The message means that the system was unable to contact its primary DNS servers, which can happen at boot time if the DNS subsystem comes online before network connectivity is established. If this message appears at other times, it could indicate network issues or that the DNS configuration is not pointing to a valid server.

Clearing the DNS Cache

The **Clear Cache** button from the GUI, or the d nsflush command (for more information about the dnsflush command, see the IronPort AsyncOS CLI Reference Guide, available from the location specified in Documentation , on page 405), clears all information in the DNS cache. You may choose to use this feature when changes have been made to your local DNS system. The command takes place immediately and may cause a temporary performance degradation while the cache is repopulated.

Configuring DNS Settings via the Graphical User Interface

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Network > DNS** page and click the **Edit Settings** button.

Step 3 Select whether to use the Internet's root DNS servers or your own internal DNS server(s), and specify authoritative DNS servers.

Step 4 If you want to use your own DNS server(s) or specify authoritative DNS servers, enter the server ID and click **Add Row.** Repeat this for each server. When entering your own DNS servers, specify a priority as well. For more information, see Specifying DNS Servers, on page 321.

Step 5 Choose an interface for DNS traffic.

Step 6 Enter the number of seconds to wait before canceling a reverse DNS lookup.

Step 7 Optionally, clear the DNS cache by clicking **Clear Cache.**

Step 8 Submit and commit your changes.

Configuring TCP/IP Traffic Routes

Some network environments require the use of traffic routes other than the standard default gateway. You can manage static routes in the GUI through the **Management Appliance > Network > Routing** page, or in the CLI by using the routeconfig command.

Managing Static Routes in the GUI

You can create, edit, or delete static routes by using the Management Appliance > Network > Routing page. You can also modify the default gateway from this page.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 On the **Management Appliance** > **Network** > **Routing** page, click **Add Route** in the route listing. The Enter a name for the route.

Step 3 Enter the destination IP address.

Step 4 Enter the gateway IP address.

Step 5 Submit and commit your changes.

Modifying the Default Gateway (GUI)

Step 1 Click Default Route in the route listing on the Routing page.
Step 2 Change the gateway IP address.
Step 3 Submit and commit your changes.

Configuring the Default Gateway

You can configure the default gateway via the GUI through the Management Appliance > Network > Routing page (see Modifying the Default Gateway (GUI), on page 324) or via the setgateway command in the CLI.

Specifying a Secure Communication Protocol

- SSL v3 is not secure and you should not use it.
- You can choose the communication protocol to be used for each of the following:
  - Updater server
  - End-user access to the spam quarantine
  - Web-based administrative interface to the appliance
  - LDAPS
- To view the currently selected protocols and available options, or to change protocols, use the sslconfig command in the command-line interface.
- Cisco update servers do not support SSL v3.
- If you are using a local (remote) update server, and for all other services and web browsers, the protocol you choose must be supported by and enabled on the server and tools you are using.
- One of the available options must be enabled for each service you use.
- Changes made using the sslconfig command require a Commit.
- Affected services will be briefly interrupted after you commit changes made using the sslconfig command.

Configuring the System Time

Note When gathering data for reports, the Security Management appliance applies the time stamp from the information that was set when you configured the time settings on the Security Management appliance. For information, see How the Security Management Appliance Gathers Data for Reports, on page 24.

To set time-related settings using the command-line interface, use the ntpconfig , settime , and settz commands.

To	Do This
Set the system time	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface. **Select Management Appliance > System Administration > Time Settings** See also Using a Network Time Protocol (NTP) Server, on page 325
Set the time zone	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface. Select **Management Appliance > System Administration > Time Zone** See also: • Selecting a GMT Offset, on page 326 • Updating Time Zone Files, on page 326

Using a Network Time Protocol (NTP) Server

You can use a Network Time Protocol (NTP) server to synchronize the Security Management appliance system clock with other computers on your network or the internet.

The default NTP server is time.sco.cisco.com .

If you will use an external NTP server, including the default NTP server, open the required port through the firewall. See Firewall Information, on page 401

Related Topics

- Configuring the System Time, on page 324
- Manually Updating Time Zone Files , on page 327

Selecting a GMT Offset

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > Time Zone**.

Step 3 Click **Edit Settings**.

Step 4 Select GMT Offset from the list of regions. The Time Zone Setting page is updated to include GMT offsets in the Time Zone field.

Step 5 Select an offset in the Time Zone field. The offset refers to the number of hours that you add or subtract to or from Greenwich Mean Time (GMT) — the local time at the prime meridian. Hours preceded by a minus sign ("-") are west of the prime meridian. A plus sign ("+") indicates locations east of the prime meridian.

Step 6 Submit and commit your changes.

Updating Time Zone Files

Whenever there is a change in the time zone rules for any country, Time Zone files on the appliance must be updated.

Automatically Updating Time Zone Files

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > Update Settings**.

Step 3 Select the **Enable automatic updates for Time zone rules** check box.

Step 4 Enter an interval. Click the **?** help on the page for important information.

Step 5 Submit and commit your changes.

Manually Updating Time Zone Files

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > Time Settings**.

Step 3 Look at the **Time Zone File Updates** section.

Step 4 If there is an available time zone file update, click **Update Now**.

Configuration File Page

For Information About This Section	See
Saving the current configuration	Saving and Importing Configuration Settings , on page 327
Loading a saved configuration	Saving and Importing Configuration Settings , on page 327
End-User Safelist/Blocklist Database (Spam Quarantine)	Backing Up and Restoring the Safelist/Blocklist , on page 185
Reset Configuration	Resetting the Configuration to Factory Defaults, on page 285

Saving and Importing Configuration Settings

Note The configuration file described in this section is used to configure Security Management appliances.

Most configuration settings for the Security Management appliance can be managed in a single configuration file. The file is maintained in Extensible Markup Language (XML) format.

You can use this file in several ways:

- In case of unexpected disaster to your primary Security Management appliance, you can quickly configure a second Security Management appliance to restore service.
- You can save the configuration file to a different system to back up and preserve crucial configuration data. If you make a mistake while configuring your appliance, you can “roll back” to the most recently saved configuration file.

- You can download the existing configuration file to view the entire configuration for an appliance quickly. (Many newer browsers include the ability to render XML files directly.) This may help you troubleshoot minor errors (like typographic errors) that may exist in the current configuration.
- You can download an existing configuration file, make changes to it, and upload it to the same appliance. This, in effect, "bypasses" both the CLI and the GUI for making configuration changes.
- You can upload an entire configuration file through FTP, or you can paste portions of a configuration file directly into the CLI.
- Because the file is in XML format, an associated document type definition (DTD) that describes all of the XML entities in the configuration file is also provided. You can download the DTD to validate an XML configuration file before uploading it. (XML validation tools are readily available on the Internet.)
- You can use the configuration file to speed configuration of another appliance, for example a cloned virtual appliance.

Managing Configuration Files

Saving and Exporting the Current Configuration File

Using the Current Configuration section of the Management Appliance > System Administration > Configuration File page, you can save the current configuration file to your local machine, save it on the appliance (placed in the configuration directory in the FTP/SCP root), or email it to the address specified.

Masking the password

Optionally, mask the user's passwords by selecting the check box. Masking a password causes the original, encrypted password to be replaced with "*****" in the exported or saved file.

Note Configuration files with masked passwords cannot be loaded back into AsyncOS.

Loading a Configuration File

The configuration file must have been saved from an appliance running the same AsyncOS version as the appliance on which you will load the configuration.

Configuration files with masked passwords cannot be loaded.

Regardless of the method, you must include the following tags at the top of your configuration:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<!DOCTYPE config SYSTEM "config.dtd">
<config>
 ... your configuration information in valid XML
</config>
```

The closing </config> tag should follow your configuration information. The values in XML syntax are parsed and validated against the DTD located in the configuration directory on your Cisco Content Security appliance. The DTD file is named config.dtd . If validation errors are reported at the command line when you use the loadconfig command, the changes are not loaded. You can download the DTD to validate configuration files outside of the appliance before uploading them.

In any import method, you can import an entire configuration file (the information defined between the highest level tags: <config></config>), or a *complete* and *unique* subsection of the configuration file, as long as it contains the declaration tags (above) and is contained within the <config></config> tags.

"Complete" means that the entire start and end tags for a given subsection as defined by the DTD are included. For example, uploading or pasting the following code causes validation errors:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<!DOCTYPE config SYSTEM "config.dtd">
<config>
  <autosupport_enabled>0</autosu
</config>
```

However, uploading or pasting the following code does not cause validation errors:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<!DOCTYPE config SYSTEM "config.dtd">
<config>
  <autosupport_enabled>0</autosupport_enabled>
</config>
```

"Unique" means that the subsection of the configuration file being uploaded or pasted is not ambiguous for the configuration. For example, a system can have only one hostname, so uploading the following code (including the declarations and <config></config> tags) is allowed:

```
<hostname>mail4.example.com</hostname>
```

However, a system can have multiple listeners defined, each with different Recipient Access Tables defined, so uploading only the following code is considered ambiguous:

```
<rat>
   <rat_entry>
     <rat_address>ALL</rat_address>
     <access>RELAY</access>
   </rat_entry>
 </rat>
```

Because it is ambiguous, it is not allowed, even though it is "complete" syntax.

Caution When uploading or pasting a configuration file or subsections of a configuration file, you have the potential to erase uncommitted changes that may be pending.

Empty Versus Omitted Tags

Use caution when uploading or pasting sections of configuration files. If you do not include a tag, then its value in the configuration is not modified when you load a configuration file. However, if you include an empty tag, then its configuration setting is cleared.

For example, uploading the following code removes all listeners from the system:

```
<listeners></listeners>
```

Caution When uploading or pasting subsections of a configuration file, you can disconnect yourself from the GUI or CLI and destroy large amounts of configuration data. Do not disable services with this command if you are not able to reconnect to the appliance using another protocol, the Serial interface, or the default settings on the Management port. Also, do not use this command if you are unsure of the exact configuration syntax as defined by the DTD. Always back up the configuration data before loading a new configuration file.

Note About Loading Passwords for Log Subscriptions

If you attempt to load a configuration file that contains a log subscription that requires a password (for example, one that will use FTP push), the loadconfig command does not warn you about the missing password. The FTP push fails and alerts are generated until you configure the correct password using the logconfig command.

Note About Character Set Encoding

The “encoding” attribute of the XML configuration file must be “ ISO-8859-1 ” regardless of the character set you may be using to manipulate the file offline. The encoding attribute is specified in the file whenever you issue the showconfig , saveconfig , or mailconfig command:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
```

Resetting the Current Configuration

Resetting the current configuration causes your Cisco Content Security appliance to revert settings back to the original factory defaults. Save your configuration prior to resetting it.

See Resetting the Configuration to Factory Defaults, on page 285.

Rolling Back to a Previously Committed Configuration

You can roll back the configuration to a previously-committed configuration.

Use the rollbackconfig command in the command-line interface to choose one of the ten most recent commits.

If you enter No when prompted to commit a rollback, the rollback will be committed the next time you commit changes.

Only users with Administrator access can use the rollbackconfig command.

Note No log messages or alerts will be generated when a previous configuration is restored.

Note Certain commits, such as re-allocating disk space to a size insufficient to hold existing data, could result in data loss.

CLI Commands for Configuration Files

The following commands enable you to manipulate the configuration files:

- `showconfig`

- `mailconfig`
- `saveconfig`
- `loadconfig`
- `rollbackconfig`
- `resetconfig` (see Resetting the Configuration to Factory Defaults, on page 285)
- `publishconfig`
- `backupconfig` (see Backing Up Security Management Appliance Data , on page 290)

The showconfig, mailconfig, and saveconfig Commands

For the configuration commands `showconfig` , `mailconfig` , and `saveconfig` , you are prompted to choose whether to include passwords in the file that will be mailed or displayed. Choosing not to include passwords leaves any password field blank. You can choose not to include passwords if you are concerned about security breaches. However, configuration files without passwords fail when loaded using the loadconfig command. See Note About Loading Passwords for Log Subscriptions, on page 330.

Note When saving, showing, or mailing your configuration file if you choose to include passwords (answer yes to “Do you want to include passwords?”), the passwords are encrypted. However, the private keys and certificates are included in unencrypted PEM format.

The `showconfig` command prints the current configuration to the screen.

```
mail3.example.com> showconfig
Do you want to include passwords? Please be aware that a configuration without passwords
will fail when reloaded with loadconfig.
<?xml version="1.0" encoding="ISO-8859-1"?>
<!DOCTYPE config SYSTEM "config.dtd">
<!--
  Product: model number
Messaging Gateway Appliance(tm)
  Model Number: model number
  Version: version of AsyncOS installed
  Serial Number: serial number
  Current Time: current time and date
[The remainder of the configuration file is printed to the screen.]
```

Use the `mailconfig` command to email the current configuration to a user. A configuration file in XML format named config.xml will be attached to the message.

```
mail3.example.com> mailconfig
Please enter the email address to which you want to send
the configuration file.
[]> administrator@example.com
Do you want to include passwords? Please be aware that a configuration without passwords
will fail when reloaded with loadconfig.  [N]> y
The configuration file has been sent to administrator@example.com.
```

The `saveconfig` command on the Security Management appliance stores and saves all of the configuration master files (ESA) with a unique filename to the configuration directory.

```
mail3.example.com> saveconfig
Do you want to include passwords? Please be aware that a configuration without passwords
will fail when reloaded with loadconfig.  [N]> y
The file C650-00065B8FCEAB-31PM121-20030630T130433.xml has been saved in the configuration
```

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

The loadconfig Command

Use the `loadconfig` command to load new configuration information into the appliance. You can load information using one of two methods:

- Placing information in the configuration directory and uploading it
- Pasting configuration information directly into the CLI

See Loading a Configuration File, on page 328 for more information.

The rollbackconfig Command

See Rolling Back to a Previously Committed Configuration , on page 330.

Uploading Configuration Changes Using the CLI

Step 1 Outside of the CLI, ensure that you are able to access the configuration directory of the appliance. See IP Interfaces and Accessing the Appliance, on page 389 for more information.

Step 2 Place an entire configuration file or subsection of a configuration file in the configuration directory of the appliance, or edit an existing configuration that was created from the `saveconfig` command.

Step 3 Within the CLI, use the `loadconfig` command to load the configuration file you placed in the directory from Step 2, or paste the text (XML syntax) directly into the CLI.
In this example, a file named changed.config.xml is uploaded and the changes are committed:

Example:

```
mail3.example.com>
l
oadconfig
1. Paste via CLI
2. Load from file
[1]> 2
Enter the name of the file to import:
[]> changed.config.xml
Values have been loaded.
Be sure to run "commit" to make these settings active.
mail3.example.com> commit
```

In this example, a new configuration file is pasted directly at the command line. (Remember to press Ctrl-D on a blank line to end the paste command.) Then the System Setup Wizard is used to change the default hostname, IP address, and gateway information. (For more information, see Running the System Setup Wizard, on page 16.) Finally, the changes are committed.

Example:

```
mail3.example.com> loadconfig
1. Paste via CLI
2. Load from file
[1]> 1
Paste the configuration file now. Press CTRL-D on a blank line when done.
```

```
[The configuration file is pasted until the end tag
</config>
. Control-D is entered on a separate line.]
Values have been loaded.
Be sure to run "commit" to make these settings active.
mail3.example.com> commit
Please enter some comments describing your changes:
[]> pasted new configuration file and changed default settings
```

Managing Disk Space

You can allocate available disk space among the features that your organization uses, up to the maximum available.

(Virtual Appliances Only) Increasing Available Disk Space

For virtual appliances running ESXi 5.5 and VMFS 5, you can allocate more than 2TB of disk space. For appliances running ESXi 5.1, the limit is 2 TB.

Note Disk space reduction in ESXi is not supported. See the VMWare documentation for information.

To add disk space to the virtual appliance instance:

Before You Begin

Carefully determine the disk space increase needed.

Step 1 Bring down the Cisco Content Security Management appliance instance.

Step 2 Increase disk space using utilities or administrative tools provided by VMWare.
See information about changing the virtual disk configuration in the VMWare documentation.

Information for ESXi 5.5 is available here:
http://pubs.vmware.com/vsphere-55/index.jsp?topic=/com.vmware.vsphere.hostclient.doc/GUID-81629CAB-72FA-42F0-9F86-F8FD0DE39E57.html

Step 3 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 4 Go to **Management Appliance > System Administration > Disk Management**and verify that your change has taken effect.

Viewing Disk Space, Quotas and Usage

To	Do This
View the total disk space available on the appliance	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface. Select **Management Appliance > System Administration > Disk Management**. Look at the values shown for "Total Space Allocated" - for example, 184G of 204G.
View the amount of disk space allocated to and currently used by each of the Security Management appliance's monitoring services	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface. Select **Management Appliance > System Administration > Disk Management**.
View the percentage of the quotas for quarantines that are currently used	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface. Select **Management Appliance > Centralized Services > System Status** and look at the Centralized Services section.

About Disk Space Maximums and Allocations

Note Centralized Reporting Disk Space on Security Management appliances is used for Email data. If you enable Centralized Email Reporting, all of the space is dedicated to the enabled feature. If you enable Email reporting data share the space and space is allocated on a first-come basis.

- Before reducing the Miscellaneous quota below current usage levels, you should delete unneeded data. See Managing Disk Space for the Miscellaneous Quota , on page 335.
- For more information about how disk space is managed for policy, virus, and outbreak quarantines, see Disk Space Allocation for Policy, Virus, and Outbreak Quarantines , on page 206 and Retention Time for Messages in Quarantines , on page 206.
- For all other data types, if you reduce the existing allocation below current usage, then the oldest data is deleted until all data fits within the new allocation amount.
- If the new quota is larger than the currently used disk space, you will not lose data.
- If you set the allocation to zero, no data is retained.

Ensuring That You Receive Alerts About Disk Space

You will begin to receive system alerts at warning level when Miscellaneous disk usage reaches 75% of the quota. You should take action when you receive these alerts.

To ensure that you receive these alerts, see Managing Alerts, on page 312.

Managing Disk Space for the Miscellaneous Quota

The Miscellaneous quota includes System data and User data. You cannot delete System data. User data that you can manage includes the following types of files:

To Manage	Do this
Log files	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface. Go to **Management Appliance > System Administration > Log Subscriptions** and: • Click the Size column heading to see which logs consume the most disk space. • Verify that you need all of the log subscriptions that are being generated. • Verify that the log level is no more verbose than necessary. • If feasible, reduce the rollover file size.
Packet captures	Go to **Help and Support**(near the upper right side of your screen) > **Packet Capture**. Delete any unneeded captures.
Configuration files (These files are unlikely to consume much disk space.)	FTP to the /data/pub directory on the appliance. To configure FTP access to the appliance, see Accessing the Appliance via FTP, on page 391
Quota size	[New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface. Go to **System Administration > Disk Management.**

Reallocating Disk Space Quotas

If disk space is allocated to features you do not use, or if the appliance frequently runs out of disk space for a particular feature and has excess space for other features, you can reallocate disk space.

If you require more space for all features, consider upgrading your hardware or allocating more disk space to your virtual appliance. See (Virtual Appliances Only) Increasing Available Disk Space , on page 333.

Before You Begin

- Changing disk allocations may impact existing data or feature availability. See information at About Disk Space Maximums and Allocations, on page 334.
- You can temporarily create space in a quarantine by manually releasing or deleting messages from the quarantine.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > Disk Management**

Step 3 Click **Edit Disk Quotas.**

Step 4 On the **Edit Disk Quotas** page, enter the amount of disk space (in gigabytes) allocated to each service.

Step 5 Click **Submit.**

Step 6 In the confirmation dialog box, click **Set New Quotas.**

Step 7 Click **Commit** to commit your changes.

Adjusting the Reference Threshold in System Health Graphs for Email Security Appliances

Note To receive alerts related to these thresholds, configure the thresholds on each managed Email Security appliance. For information, see information about configuring thresholds for system health in the user guide or online help for your Email Security appliance release. You can also run on-demand system health checks from individual appliances. See information about checking the health of your appliance in the user guide or online help for your Email Security appliance release.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Click **Management Appliance > System Administration > System Health**.

Step 3 Click **Edit Settings**.

Step 4 Configure options.

Option	Description
Overall CPU Usage	Default: 85%
Memory Page Swapping	Default: 5000 pages
Maximum Messages in Work Queue	Default: 500 messages

Step 5 Submit and commit your changes.

SSO Using SAML 2.0

About SSO and SAML 2.0

Cisco Content Security Management appliance now supports SAML 2.0 SSO so that the end users can access the Spam Quarantine using the same credentials that are used to access other SAML 2.0 SSO enabled services within their organization. For instance, you have enabled Ping Identity as your SAML identity provider (IdP) and has accounts on Rally, Salesforce, and Dropbox which have been SAML 2.0 SSO enabled. When you configure Cisco Content Security Management appliance to support SAML 2.0 SSO as a Service Provider (SP), end users will be able to sign in once and have access to all these services including Spam Quarantine.

SAML 2.0 SSO Workflow

The SAML 2.0 SSO workflow is displayed in the following figure:

Figure 19: SAML 2.0 SSO Workflow

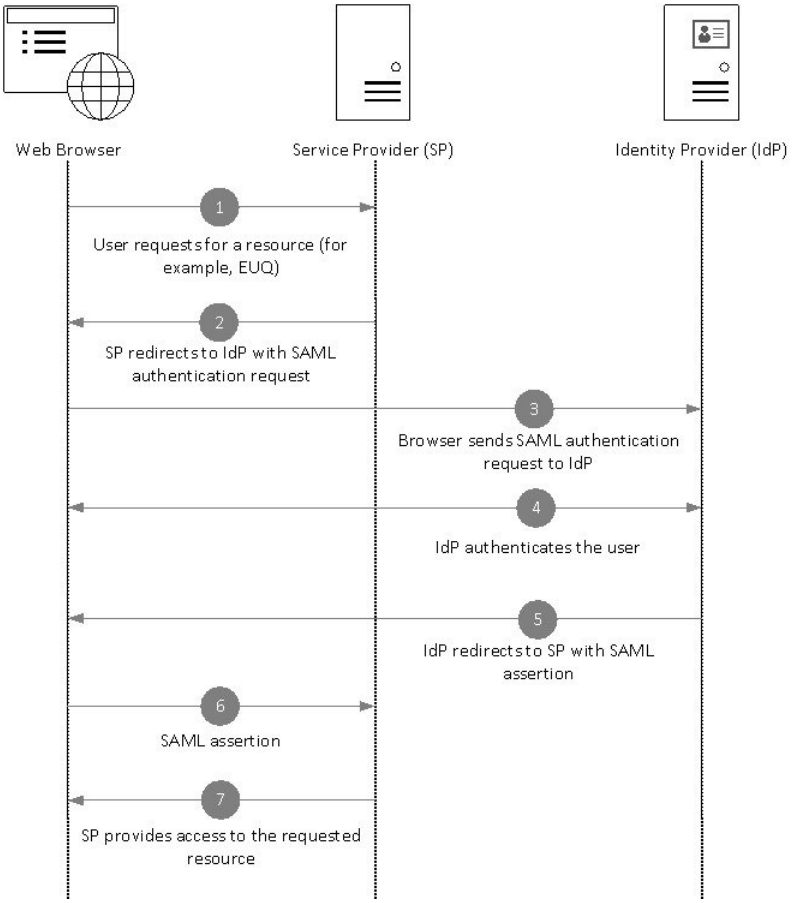

Workflow

1. The end user uses a web browser to request a resource from the service provider (your appliance). For example, the end user clicks on the spam quarantine link in a spam notification.
2. The service provider redirects the request to the web browser with SAML authentication request.
3. The web browser relays the SAML authentication request to the identity provider.
4. The identity provider authenticates the end user. The identity provider displays a login page to the end user and the end user logs in.
5. The identity provider generates the SAML assertion and sends it back to the web browser.
6. The web browser relays the SAML assertion to the service provider.
7. The service provider grants access to the requested resource.

Guidelines and Limitations for SAML 2.0

Logout

When end users log out of Spam Quarantine, they are not logged out of other SAML 2.0 SSO enabled applications.

General

You can configure only one instance of service provider and identity provider on Cisco Content Security Management appliance.

Spam Quarantine Access for Administrators

If you are enabling SSO for Spam Quarantine, keep in mind that the administrators will no longer be able to access the Spam Quarantine using the Spam Quarantine URL (http://<appliance_hostname>:<port>). Administrators can access the Spam Quarantine using the web interface (**Email > Message Quarantine > Spam Quarantine**).

How to Configure SSO for Spam Quarantine

	Do This	More Info
Step 1	Review the prerequisites.	Prerequisites, on page 340
Step 2	Configure your appliance as a service provider.	Configure Cisco Content Security Management Appliance as a Service Provider, on page 340
Step 3	[On IDP] Configure the identity provider to work with your appliance.	Configure the Identity Provider to Communicate with Cisco Content Security Management Appliance, on page 342
Step 4	Configure identity provider settings on your appliance.	Configure Identity Provider Settings on Cisco Content Security Management Appliance, on page 343
Step 5	Enable SSO for Spam Quarantine on your appliance.	Enable SSO for Spam Quarantine, on page 344
Step 6	Notify the end users about the new authentication mechanism.	

Prerequisites

- Verify whether the identity provider used by your organization is supported by Cisco Content Security Management Appliance. The following are the supported identity providers:
  - Microsoft Active Directory Federation Services (AD FS) 2.0
  - Ping Identity PingFederate 7.2
- Obtain the following certificates that are required to secure the communication between your appliance and the identity provider:
  - If you want your appliance to sign SAML authentication requests or if want your identity provider to encrypt SAML assertions, obtain a self signed certificate or a certificate from a trusted CA and the associated private key.
  - If you want the identity provider to sign SAML assertions, obtain the identity provider's certificate. Your appliance will use this certificate to verify the signed SAML assertions.

Configure Cisco Content Security Management Appliance as a Service Provider

Before You Begin

Review the Prerequisites, on page 340

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > System Administration > SAML**.

Step 3 Under the Service Provider section, click **Add Service Provider**.

Step 4 Enter the following details:

Field	Description
Profile Name	Enter a name for the service provider profile.
Configuration Settings	
Entity ID	Enter a globally unique name for the service provider (in this case, your appliance). The format of the service provider Entity ID is typically a URI.
Name ID Format	The format that the identity provider should use to specify the user in the SAML assertion. This field is not configurable. You will need this value while configuring the identity provider.

Field	Description
Assertion Consumer URL	The URL to which the identity provider should send the SAML assertion after authentication has successfully completed. In this case, this is the URL to your spam quarantine. This field is not configurable. You will need this value while configuring the identity provider.
SP Certificate	**Note** The private key must be in .pem format. **Signing Authentication Requests** If you want the appliance to sign the SAML authentication requests: 1 Upload the certificate and the associated private key. 2 Enter the passphrase for the private key. 3 Select **Sign Request**. **Decrypt Encrypted Assertions** If you plan to configure your identity provider to encrypt SAML assertions: 1 Upload the certificate and the associated private key. 2 Enter the passphrase for the private key.
Sign Assertions	If you want the identity provider to sign the SAML assertions, select **Sign Assertions**. If you select this option, you must add the identity provider's certificate to the appliance. See Configure Identity Provider Settings on Cisco Content Security Management Appliance, on page 343.
Organization Details	Enter the details of your organization. Identity provider uses this information in the error logs.
Technical Contact	Enter the email address of the technical contact. Identity provider uses this information in the error logs.

Step 5 Click **Submit**.

Step 6 Note down the service provider metadata (Entity ID and Assertion Customer URL) displayed on the SSO Settings page and the Name ID Format displayed on the Service Provider Settings page. You will need these details while configuring the service provider settings on the identity provider.
Optionally, you can export the metadata as a file. Click **Export Metadata** and save the metadata file. Some identity providers allow you to load service provider details from a metadata file.

What to Do Next

Configure the identity provider to communicate with your appliance. See Configure the Identity Provider to Communicate with Cisco Content Security Management Appliance, on page 342

Configure the Identity Provider to Communicate with Cisco Content Security Management Appliance

Before You Begin

Make sure that you have:

- Configured your appliance as a service provider. See Configure Cisco Content Security Management Appliance as a Service Provider, on page 340.
- Copied the service provider metadata details or exported the metadata file. See Configure Cisco Content Security Management Appliance as a Service Provider, on page 340.

Step 1 On the identity provider, do one of the following:

- Manually configure the details of the service provider (your appliance).
- If your identity provider allows you to load the service provider details from a metadata file, import the metadata file.

If you have configured your appliance to sign the SAML authentication requests or you plan to encrypt SAML assertions, make sure that you add the relevant certificate to the identity provider.

For identity provider-specific instructions, see:

- Configure AD FS 2.0 to Communicate with Cisco Content Security Management Appliance, on page 342
- Configure PingFederate 7.2 to Communicate with Cisco Content Security Management Appliance, on page 343
- **Configuring the Appliance as an Identity Provider** section in the *User Guide for AsyncOS for Cisco Web Security Appliances* http://www.cisco.com/c/en/us/support/security/web-security-appliance/products-user-guide-list.html

Step 2 Note down the identity provider metadata or export the metadata as a file.

What to Do Next

Configure the identity provider settings on your appliance. See Configure Identity Provider Settings on Cisco Content Security Management Appliance, on page 343.

Configure AD FS 2.0 to Communicate with Cisco Content Security Management Appliance

The following are the high level tasks you need to perform to configure AD FS 2.0 to communicate with your appliance. For complete and detailed instructions, see Microsoft documentation.

- Add the service provider's (appliance's) Assertion Consumer URL as a relaying party.
- Enter the service provider's (appliance's) Entity ID under Relaying Party Trusts > Properties > Identifiers > Relaying Party Identifier. Make sure that this value is same as the Entity ID value in the Service Provider settings on your appliance.

- If you have configured your service provider (appliance) to send signed SAML authentication requests, upload the service provider's certificate (used to sign authentication requests) in .cer format under Relaying Party Trusts > Properties > Signature.
- If you plan to configure AD FS to send encrypted SAML assertions, upload the service provider's (appliance's) certificate in .cer format under Relaying Party Trusts > Properties > Encryption.
- Set the Secure-hash Algorithm to SHA-1 under Relaying Party Trusts > Properties > Advanced.
- Edit the Claim Rule and add an Issuance Transform Rule to send the LDAP attribute for email address as an outgoing claim type (email address).
- Add a custom rule to include SPNameQualifier in the response. The following is a sample custom rule:

```
c:[Type == "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailaddress"] =>
issue(Type = "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/nameidentifier", Issuer
 =
 c.Issuer, OriginalIssuer = c.OriginalIssuer, Value = c.Value, ValueType = c.ValueType,
Properties["http://schemas.xmlsoap.org/ws/2005/05/identity/claimproperties/format"] =
"urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress", Properties
["http://schemas.xmlsoap.org/ws/2005/05/identity/claimproperties/spnamequalifier"] =
"https://<appliance-hostname>:83");
```

Configure PingFederate 7.2 to Communicate with Cisco Content Security Management Appliance

The following are the high level tasks you need to perform to configure PingFederate 7.2 to communicate with your appliance. For complete and detailed instructions, see Ping Identity documentation.

- Add your service provider's (appliance's) Assertion Consumer URL as an endpoint under protocol settings.
- Enter the service provider's (appliance's) Entity ID under SP Connection > General Info > Partner's Entity ID (Connection ID). Make sure that this value is same as the Entity ID value in the Service Provider settings on your appliance.
- If you have configured your service provider (appliance) to send signed SAML authentication requests, upload the service provider's certificate under Signature Verification section (SP Connection > Credentials > Signature Verification > Signature Verification Certificate).
- If you plan to configure PingFederate to send encrypted SAML assertions, upload the service provider's (appliance's) certificate under Signature Verification section (SP Connection > Credentials > Signature Verification > Select XML Encryption Certificate).
- Edit Attribute Contract to send the LDAP attribute- email address (Attribute Sources & User Lookup > Attribute Contract Fulfillment).

Configure Identity Provider Settings on Cisco Content Security Management Appliance

Before You Begin

Make sure that you have:

- Configured the identity provider to communicate with your appliance. See Configure the Identity Provider to Communicate with Cisco Content Security Management Appliance, on page 342.

- Copied the identity provider metadata details or the exported metadata file.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance** > **System Administration** > **SAML**.

Step 3 Under the Identity Provider section, click **Add Identity Provider**.

Step 4 Enter the following details:

Field	Description
Profile Name	Enter a name for the identity provider profile.
Configuration Settings (Manually Configure Identity Provider Settings)	
Entity ID	Enter a globally unique name for the identity provider. The format of the identity provider Entity ID is typically a URI.
SSO URL	Specify the URL to which the service provider must send the SAML authentication requests.
Certificate	If the identity provider signs the SAML assertion, you must upload the identity provider's signing certificate.
Configuration Settings (Importing Identity Provider Metadata)	
Import IDP Metadata	Click **Import Metadata** and select the metadata file.

Step 5 Submit and commit your changes.

What to Do Next

Enable SSO for Spam Quarantine, on page 344

Enable SSO for Spam Quarantine

Before You Begin

Make sure that you have:

- Configured all the settings on **Management Appliance** > **System Administration** > **SAML** page.

- Enabled Spam Quarantine. See Spam Quarantine, on page 173.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance** > **Centralized Services** > **Spam Quarantine**.

Step 3 Click **Edit Settings** and scroll down to the End-User Quarantine Access section.

Step 4 Make sure that you have enabled End-User Quarantine Access.

Step 5 Set the End-User Authentication method to **SAML2.0**.

Step 6 (Optional) Specify whether to display message bodies before messages are released.

Step 7 Submit and commit your changes.

What to Do Next

Notify the end users about the new authentication mechanism.

Customizing Your View

Using Favorite Pages

(Locally-authenticated administrative users only.) You can create a quick-access list of the pages you use most.

To	Do This
Add pages to your favorites list	Navigate to the page to add, then choose **Add This Page To My Favorites**from the My Favorites menu near the top right corner of the window. No commit is necessary for changes to My Favorites.
Reorder favorites	Choose **My Favorites > View All My Favorites** and drag favorites into the desired order.
Edit favorite page, name, or description	Choose **My Favorites > View All My Favorites** and click the name of the favorite to edit.
Delete favorites	Choose **My Favorites > View All My Favorites** and delete favorites.
Go to a favorite page	Choose a page from the **My Favorites** menu near the top right corner of the window.
Return to the main interface	Choose any favorite, or click the **Return to previous page** at the bottom of the page.

Setting Preferences

Administrative users configured on the Security Management appliance

Locally-authenticated users can choose the following preferences, which apply each time the user logs in to the Security Management appliance:

- Language (applies to the GUI)
- Landing page (the page displayed after login)
- Default time range for report pages (available options are a subset of the time ranges available for Email reporting pages)
- Number of rows visible in tables on report pages

Exact options depend on the user role.

To set these preferences, choose **Options > Preferences**. (The Options menu is at the top right side of the GUI window.) Submit your changes when done. Commit is not required.

Tip To return to the page you were viewing before you accessed the Preferences page, click the **Return to previous page** link at the bottom of the page.

Externally authenticated users

Externally authenticated users can choose the display language directly in the Options menu.

Improving Web Interface Rendering

For better web interface rendering, Cisco recommends that you enable Internet Explorer Compatibility Mode Override.

Note If enabling this feature is against your organizational policy, you may disable this feature.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance** > **System Administration** > **General Settings**.

Step 3 Select **Override IE Compatibility Mode** check box.

Step 4 Submit and commit your changes.

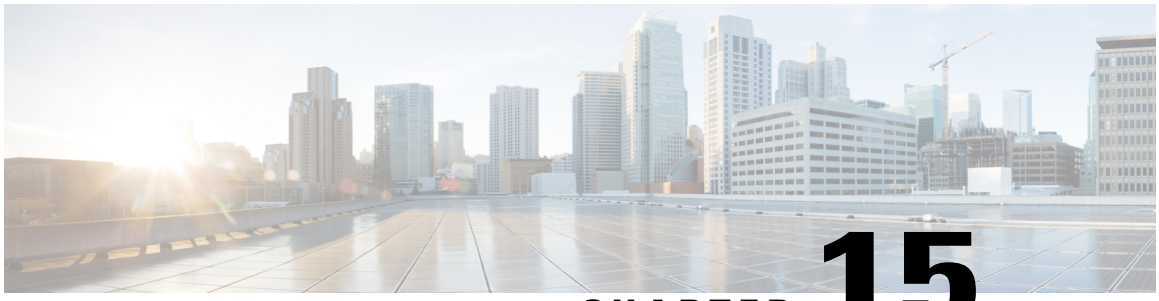

CHAPTER 15

Logging

This chapter contains the following sections:

Logging Overview

Log files record regular operations, as well as exceptions, for activity on the system. Use the logs for monitoring the Cisco Content Security appliance, troubleshooting, and evaluating system performance.

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

Logging Versus Reporting

Use logging data to debug message flow, reveal basic day-to-day operational information such as FTP connection details, HTTP log files, and for compliance archiving.

You can access this logging data directly on the Email Security appliance or send it to any external FTP server for archival or reading. You can either FTP to the appliance to access the logs or push the plain text logs to an external server for backup purposes.

To view reporting data, use the Report pages on the appliance GUI. You cannot access the underlying data in any way, and this data cannot be sent to anything but a Cisco Content Security Management appliance.

Note The Security Management appliance pulls information for all reporting and tracking with the exception of spam quarantine data. This data is pushed from the ESA.

Log Retrieval

Log files can be retrieved with the file transfer protocols described in Table 15-1. You set the protocol when you create or edit a log subscription in the GUI, or by using the logconfig command in the CLI.

Table 58:

FTP Poll	With this type of file transfer, a remote FTP client accesses the appliance to retrieve log files by using the user name and password of an administrator-level or operator-level user. When configuring a log subscription to use the FTP poll method, you must supply the maximum number of log files to retain. When the maximum number is reached, the system deletes the oldest file.
FTP Push	With this type of file transfer, the Cisco Content Security appliance periodically pushes log files to an FTP server on a remote computer. The subscription requires a user name, password, and destination directory on the remote computer. Log files are transferred based on the configured rollover schedule.
SCP Push	With this type of file transfer, the Cisco Content Security appliance periodically pushes log files to an SCP server on a remote computer. This method requires an SSH SCP server on a remote computer using the SSH2 protocol. The subscription requires a user name, SSH key, and destination directory on the remote computer. Log files are transferred based on the configured rollover schedule.
Syslog Push	With this type of file transfer, the Cisco Content Security appliance sends log messages to a remote syslog server. This method conforms to RFC 3164. You must submit a hostname for the syslog server and use either UDP or TCP for log transmission. The port used is 514. A facility can be selected for the log; however, a default for the log type is preselected in the drop-down menu. Only text-based logs can be transferred using syslog push.

Filename and Directory Structure

AsyncOS creates a directory for each log subscription based on the log name specified in the log subscription. The filenames of logs in the directory consist of the filename specified in the log subscription, the timestamp when the log file was started, and a single-character status code. The following example shows the convention for the directory and filename:

/<Log_Name>/<Log_Filename>.@<timestamp>.<statuscode>

Status codes may be .c (signifying “current”) or .s (signifying “saved”). You should only transfer log files with the saved status.

Log Rollover and Transfer Schedule

When you create a log subscription, you specify the trigger(s) for when the logs roll over, the old file is transferred, and a new log file is created.

Choose between the following triggers:

- File size
- Time
  - At a specified interval (in seconds, minutes, hours, or days)

    Follow the example on the screen when entering values.

    To enter a composite interval, such as two-and-a-half hours, follow the example 2h30m .

    or
  - Every day, at the time(s) you specify

    or
  - On the days of the week that you select, at the time(s) you specify

When you specify times, use the 24-hour format, for example 23:00 for 11pm.

To schedule multiple rollover times in a day, separate times with a comma. For example, to roll over logs at midnight and noon, enter 00:00, 12:00

Use an asterisk (*) as a wildcard. For example, to roll over logs exactly at every hour and half-hour, enter *:00, *:30

When the specified limit is reached (or the first limit is reached, if you have configured both size- and time-based limits), the log file is rolled over. Log subscriptions based on the FTP poll transfer mechanism create files and store them in the FTP directory on the appliance until they are retrieved or until the system needs more space for log files.

Note If a rollover is in progress when the next limit is reached, the new rollover is skipped. An error will be logged and an alert sent.

Timestamps in Log Files

The following log files include the beginning and ending date of the log itself, the version of AsyncOS, and the GMT offset (provided in seconds at the beginning of the log):

- Mail log
- Safelist/blocklist log
- System log

Logs Enabled by Default

The Security Management appliance is preconfigured with the following log subscriptions enabled.

Table 59: Preconfigured Log Subscriptions

Log Name	Log Type	Retrieval Method
cli_logs	CLI audit logs	FTP Poll
euq_logs	Spam quarantine logs	FTP Poll
euqgui_logs	Spam quarantine GUI logs	FTP Poll
gui_logs	HTTP logs	FTP Poll
mail_logs	Text mail logs	FTP Poll
reportd_logs	Reporting logs	FTP Poll
reportqueryd_logs	Reporting query logs	FTP Poll
slbld_logs	Safelist/blocklist logs	FTP Poll
smad_logs	SMA logs	FTP Poll
system_logs	System logs	FTP Poll
trackerd_logs	Tracking logs	FTP Poll

All preconfigured log subscriptions have the logging level set to Information. For more information about log levels, see Setting the Log Level, on page 372.

You can configure additional log subscriptions depending on the license keys that you have applied. For information about creating and editing log subscriptions, see Log Subscriptions, on page 371.

Log Types

Summary of Log Types

A log subscription associates a log type with a name, a logging level, and other characteristics such as file size and destination information. Multiple subscriptions for all log types, except configuration history logs, are permitted. The log type determines the data that are recorded in the log. You select the log type when you create a log subscription. See Log Subscriptions, on page 371 for more information.

AsyncOS generates the following log types:

Table 60: Log Types

Log Type	Description
Authentication Logs	The authentication log records successful logins and unsuccessful login attempts, for locally and externally authenticated users, for both GUI and CLI access to the Security Management appliance. In Debug and more verbose modes, if external authentication is turned on, all LDAP queries appear in these logs.
Backup Logs	Backup logs record the backup process from start to finish. Information about backup scheduling is in the SMA logs.
CLI Audit Logs	The CLI audit logs record all CLI activity on the system.
Configuration History Logs	Configuration history logs record the following information: What changes were made on the Security Management appliance, and when were the changes made? A new configuration history log is created each time a user commits a change.
FTP Server Logs	FTP logs record information about the FTP services enabled on the interface. Connection details and user activity are recorded.
GUI logs	GUI logs include a history of page refreshes in the web interface, session data, and the pages a user accesses. You can use the gui_log to track user activity or investigate errors that users see in the GUI. The error traceback will normally be in this log. GUI logs also include information about SMTP transactions, for example information about scheduled reports emailed from the appliance.
HTTP Logs	HTTP logs record information about the HTTP and secure HTTP services enabled on the interface. Because the graphical user interface (GUI) is accessed through HTTP, the HTTP logs are essentially the GUI equivalent of the CLI audit logs. Session data (for example, new sessions and expired sessions) are recorded, as well as the pages accessed in the GUI.
Haystack logs	Haystack logs record web transaction tracking data processing.

Log Type	Description
Text Mail Logs	Text mail logs record information about the operations of the email system (for example, message receiving, message delivery attempts, opening and closing connections, bouncing messages, and so forth). For important information about when attachment names are included in mail logs, see Tracking Service Overview, on page 159.
LDAP Debug Logs	Use these logs to debug problems when you are configuring LDAP in System Administration > LDAP. For example, these logs record the results of clicking the Test Server and Test Queries buttons. For information about failed LDAP authentications, see the Authentication logs.
NTP Logs	NTP logs record the conversation between the appliance and any configured Network Time Protocol (NTP) servers. For information about configuring NTP servers, see Configuring the System Time, on page 324.
Reporting Logs	Reporting logs record actions associated with the processes of the centralized reporting service.
Reporting Query Logs	Reporting query logs record actions associated with the reporting queries that are run on the appliance.
SMA Logs	SMA logs record actions associated with general Security Management appliance processes, not including the processes of the centralized reporting, centralized tracking, and spam quarantine services. These logs include information about backup scheduling.
SNMP Logs	SNMP logs record debug messages related to the SNMP network management engine. In Trace or Debug mode, this includes SNMP requests to the Security Management appliance.
Safelist/Blocklist Logs	Safelist/blocklist logs record data about the safelist/blocklist settings and database.

Log Type	Description
Spam Quarantine GUI Logs	Spam quarantine GUI logs record actions associated with the spam quarantine GUI, such as quarantine configuration through the GUI, end user authentication, and end user actions (for example, releasing email).
Spam Quarantine Logs	Spam quarantine logs record actions associated with the spam quarantine processes.
Status Logs	Status logs record system statistics found in the CLI status commands, including status detail and dnsstatus . The period of recording is set using the setup subcommand in logconfig . Each counter or rate reported in status logs is the value since the last time the counter was reset.
System Logs	System logs record the following: boot information, DNS status information, and comments users typed using the commit command. System logs are useful for troubleshooting the state of the appliance.
Tracking Logs	Tracking logs record actions associated with the processes of the tracking service. Tracking logs are a subset of the mail logs.
Updater Logs	Information about service updates, such as time zone updates.
Upgrade Logs	Status information about upgrade download and installation.

Log Type Comparison

The following tablesummarizes the characteristics of each log type.

Table 61: Log Type Comparison

						Contains					
	Transactional	Stateless	Recorded as Text	Recorded as Binary	Header Logging	Periodic Status Information	Message Receiving Information	Delivery Information	Individual Hard Bounces	Individual Soft Bounces	Configuration Information
Authentication Logs	•		•								

						Contains					
Backup Logs	•		•								
CLI Audit Logs	•		•			•					
Configuration History Logs	•		•								•
FTP Server Logs	•		•			•					
HTTP Logs	•		•			•					
Haystack Logs	•		•								
Text Mail Logs	•		•		•	•	•	•	•	•	
LDAP Debug Logs	•		•								
NTP Logs	•		•			•					
Reporting Logs	•		•			•					
Reporting Query Logs	•		•			•					
SMA Logs	•		•			•					
SNMP Logs	•		•								
Safelist/ Blocklist Logs	•		•			•					

						Contains					
Spam Quarantine GUI	•		•			•					
Spam Quarantine	•		•			•					
Status Logs		•	•			•					
System Logs	•		•			•					
Tracking Logs	•			•	•		•	•	•	•	
Updater Logs	•		•								

Using Configuration History Logs

A configuration history log consists of a configuration file with an additional section listing the name of the user, a description of where in the configuration the user made changes, and the comment the user entered when committing the change. Each time a user commits a change, a new log is created containing the configuration file after the change.

Example

In this example, the configuration history log shows that the user (admin) added a guest user to the table that defines which local users are allowed to log in to the system.

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<!DOCTYPE config SYSTEM "config.dtd">
<!--
  XML generated by configuration change.
  Change comment: added guest user
  User: admin
  Configuration are described as:
    This table defines which local users are allowed to log into the system.
  Product: M160 Messaging Gateway(tm) Appliance
  Model Number: M160
  Version: 6.7.0-231
  Serial Number: 000000000ABC-D000000
  Number of CPUs: 1
  Memory (GB): 4
  Current Time: Thu Mar 26 05:34:36 2009
  Feature "Centralized Configuration Manager": Quantity = 10, Time Remaining = "25 days"
  Feature "Centralized Reporting": Quantity = 10, Time Remaining = "9 days"
  Feature "Centralized Tracking": Quantity = 10, Time Remaining = "30 days"
  Feature "Centralized Spam Quarantine": Quantity = 10, Time Remaining = "30 days"
  Feature "Receiving": Quantity = 1, Time Remaining = "Perpetual"
-->
<config>
```

Using CLI Audit Logs

The following table describes the statistics recorded in CLI audit logs.

Table 62: CLI Audit Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
PID	Process ID for the particular CLI session in which the command was entered.
Message	The message consists of the CLI command that was entered, the CLI output (including menus, lists, and so forth), and the prompt that appears.

Example

In this example, the CLI audit log shows that, for PID 16434, the following CLI commands were entered: who , textconfig .

```
Thu Sep  9 14:35:55 2004 Info: PID 16434: User admin entered 'who'; prompt was
'\nmail3.example.com> '
Thu Sep  9 14:37:12 2004 Info: PID 16434: User admin entered 'textconfig'; prompt was
'\nUsername  Login Time  Idle Time  Remote Host  What\n
========  ==========  =========  ===========  ====\n
admin     Wed 11AM    3m 45s     10.1.3.14    tail\nadmin     02:32PM     0s        10.1.3.14
    cli\nmail3.example.com> '
Thu Sep  9 14:37:18 2004 Info: PID 16434: User admin entered ''; prompt was '\nThere are
no text resources currently defined.\n\n\nChoose the operation you want to perform:\n- NEW
 - Create a new text resource.\n- IMPORT - Import a text resource from a file.\n[]> '
```

Using FTP Server Logs

The following table describes the statistics recorded in FTP server logs.

Table 63: FTP Server Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
ID	Connection ID. A separate ID for each FTP connection.
Message	The message section of the log entry can be logfile status information, or FTP connection information (login, upload, download, logout, and so forth).

Example

In this example, the FTP server log records a connection (ID:1). The IP address of the incoming connection is shown, as well as the activity (uploading and downloading files) and the logout.

```
Wed Sep  8 18:03:06 2004 Info: Begin Logfile
Wed Sep  8 18:03:06 2004 Info: Version: 4.0.0-206 SN: 00065BF3BA6D-9WFWC21
Wed Sep  8 18:03:06 2004 Info: Time offset from UTC: 0 seconds
Wed Sep  8 18:03:06 2004 Info: System is coming up
Fri Sep 10 08:07:32 2004 Info: Time offset from UTC: -25200 seconds
Fri Sep 10 08:07:32 2004 Info: ID:1 Connection from 10.1.3.14 on 172.19.0.86
Fri Sep 10 08:07:38 2004 Info: ID:1 User admin login SUCCESS
Fri Sep 10 08:08:46 2004 Info: ID:1 Upload wording.txt 20 bytes
Fri Sep 10 08:08:57 2004 Info: ID:1 Download words.txt 1191 bytes
Fri Sep 10 08:09:06 2004 Info: ID:1 User admin logout
```

Using HTTP Logs

The following tabledescribes the statistics recorded in HTTP logs

Table 64: Statistics Recorded in HTTP Logs

Statistic	Description
Timestamp	Time that the bytes were transmitted.
ID	Session ID.
req	IP address of machine connecting.
user	User name of user connecting.
Message	Information regarding the actions performed. May include GET or POST commands or system status, and so forth.

Example

In this example, the HTTP log shows the admin user's interaction with the GUI (for example, running the System Setup Wizard).

```
Wed Sep  8 18:17:23 2004 Info: http service on 192.168.0.1:80 redirecting to https port 443
Wed Sep  8 18:17:23 2004 Info: http service listening on 192.168.0.1:80
Wed Sep  8 18:17:23 2004 Info: https service listening on 192.168.0.1:443
Wed Sep  8 11:17:24 2004 Info: Time offset from UTC: -25200 seconds
Wed Sep  8 11:17:24 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg POST
/system_administration/system_setup_wizard HTTP/1.1 303
Wed Sep  8 11:17:25 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg GET
/system_administration/ssw_done HTTP/1.1 200
Wed Sep  8 11:18:45 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg GET
/monitor/incoming_mail_overview HTTP/1.1 200
Wed Sep  8 11:18:45 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg GET
/monitor/mail_flow_graph?injector=&width=365&interval=0&type=recipientsin&height=190 HTTP/1.1
 200
Wed Sep  8 11:18:46 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg GET
/monitor/classification_graph?injector=&width=325&interval=0&type=recipientsin&height=190
HTTP/1.1 200
Wed Sep  8 11:18:49 2004 Info: req:10.10.10.14 user:admin id:iaCkEh2h5rZknQarAecg GET
/monitor/quarantines HTTP/1.1 200
```

Using Spam Quarantine Logs

The following table describes the statistics recorded in spam quarantine logs.

Table 65: Spam Quarantine Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken (messages quarantined, released from quarantine, and so forth).

Example

In this example, the log shows two messages (MID 8298624 and MID 8298625) being released from the quarantine to admin@example.com.

```
Mon Aug 14 21:41:47 2006 Info: ISQ: Releasing MID [8298624, 8298625] for all
Mon Aug 14 21:41:47 2006 Info: ISQ: Delivering released MID 8298624 (skipping work queue)
Mon Aug 14 21:41:47 2006 Info: ISQ: Released MID 8298624 to admin@example.com
Mon Aug 14 21:41:47 2006 Info: ISQ: Delivering released MID 8298625 (skipping work queue)
Mon Aug 14 21:41:47 2006 Info: ISQ: Released MID8298625 to admin@example.com
```

Using Spam Quarantine GUI Logs

The following table shows the statistics recorded in spam quarantine GUI logs.

Table 66: Spam Quarantine GUI Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken, including user authentication, and so forth.

Example

In this example, the log shows a successful authentication, login, and logout:

Table 67: Spam Quarantine GUI Log Example

`Fri Aug 11 22:05:28 2006 Info: ISQ: Serving HTTP on 192.168.0.1, port 82`
`Fri Aug 11 22:05:29 2006 Info: ISQ: Serving HTTPS on 192.168.0.1, port 83`
`Fri Aug 11 22:08:35 2006 Info: Authentication OK, user admin`
`Fri Aug 11 22:08:35 2006 Info: logout:- user:pqufOtL6vyI5StCqhCfO session:10.251.23.228`
`Fri Aug 11 22:08:35 2006 Info: login:admin user:pqufOtL6vyI5StCqhCfO session:10.251.23.228`
`Fri Aug 11 22:08:44 2006 Info: Authentication OK, user admin`

Using Text Mail Logs

These logs contain details of email receiving, email delivery, and bounces. Status information is also written to the mail log every minute. These logs are a useful source of information to understand delivery of specific messages and to analyze system performance.

These logs do not require any special configuration. However, you must configure the system properly to view attachment names, and attachment names may not always be logged. For specifics, see Tracking Service Overview, on page 159.

The following table shows the information displayed in text mail logs.

Table 68: Text Mail Log Statistics

Statistic	Description
ICID	Injection Connection ID. This is a numerical identifier for an individual SMTP connection to the system. A single message or thousands of individual messages can be sent over one SMTP connection to the system.
DCID	Delivery Connection ID. This is a numerical identifier for an individual SMTP connection to another server, for delivery of one to thousands of messages, each with some or all of its RIDs being delivered in a single message transmission.
RCID	RPC Connection ID. This is a numerical identifier for an individual RPC connection to the spam quarantine. It is used to track messages as they are sent to and from the spam quarantine.
MID	Message ID. Use this to track messages as they flow through the logs.

Statistic	Description
RID	Recipient ID. Each message recipient is assigned an ID.
New	New connection initiated.
Start	New message started.

Sample Text Mail Log

Use the following sample as a guide to interpret log files.

Note Individual lines in log files are *not* numbered. They are numbered here only for sample purposes.

Table 69: Text Mail Log Detail

1	`Mon Apr 17 19:56:22 2003 Info: New SMTP ICID 5 interface Management (10.1.1.1) address 10.1.1.209 reverse dns host remotehost.com verified yes`
2	`Mon Apr 17 19:57:20 2003 Info: Start MID 6 ICID 5`
3	`Mon Apr 17 19:57:20 2003 Info: MID 6 ICID 5 From: <sender@remotehost.com>`
4	`Mon Apr 17 19:58:06 2003 Info: MID 6 ICID 5 RID 0 To: <mary@yourdomain.com>`
5	`Mon Apr 17 19:59:52 2003 Info: MID 6 ready 100 bytes from <sender@remotehost.com>`
6	`Mon Apr 17 19:59:59 2003 Info: ICID 5 close`
7	`Mon Mar 31 20:10:58 2003 Info: New SMTP DCID 8 interface 192.168.42.42 address 10.5.3.25`
8	`Mon Mar 31 20:10:58 2003 Info: Delivery start DCID 8 MID 6 to RID [0]`
9	`Mon Mar 31 20:10:58 2003 Info: Message done DCID 8 MID 6 to RID [0]`
10	`Mon Mar 31 20:11:03 2003 Info: DCID 8 close`

The following table can be used as a guide to reading the previous log file.

Table 70: Detail of Text Mail Log Example

Line Number	Description
1	A new connection is initiated into the system and assigned an Injection ID (ICID) of "5." The connection was received on the Management IP interface and was initiated from the remote host at 10.1.1.209.
2	The message is assigned a Message ID (MID) of "6" after the MAIL FROM command is issued from the client.
3	The sender address is identified and accepted.
4	The recipient is identified and assigned a Recipient ID (RID) of "0."
5	MID 5 is accepted, written to disk, and acknowledged.
6	Receiving is successful and the receiving connection closes.
7	The message delivery process starts. It is assigned a Delivery Connection ID (DCID) of "8" from 192.168.42.42 and to 10.5.3.25.
8	The message delivery starts to RID "0."
9	Delivery is successful for MID 6 to RID "0."
10	The delivery connection closes.

Examples of Text Mail Log Entries

The following examples show log entries based on various cases.

Message Receiving

A message is injected into the appliance for a single recipient. The message is successfully delivered.

```
Wed Jun 16 21:42:34 2004 Info: New SMTP ICID 282204970 interface mail.example.com (1.2.3.4)
 address 2.3.4.5 reverse dns host unknown verified no
Wed Jun 16 21:42:34 2004 Info: ICID 282204970 SBRS None
Wed Jun 16 21:42:35 2004 Info: Start MID 200257070 ICID 282204970
Wed Jun 16 21:42:35 2004 Info: MID 200257070 ICID 282204970 From: <someone@foo.com>
Wed Jun 16 21:42:36 2004 Info: MID 200257070 ICID 282204970 RID 0 To: <user@example.com>
Wed Jun 16 21:42:38 2004 Info: MID 200257070 Message-ID '<37gva9$5uvbhe@mail.example.com>'
Wed Jun 16 21:42:38 2004 Info: MID 200257070 Subject 'Hello'
Wed Jun 16 21:42:38 2004 Info: MID 200257070 ready 24663 bytes from <someone@foo.com>
Wed Jun 16 21:42:38 2004 Info: MID 200257070 antivirus negative
Wed Jun 16 21:42:38 2004 Info: MID 200257070 queued for delivery
```

```
Wed Jun 16 21:42:38 2004 Info: New SMTP DCID 2386069 interface 1.2.3.4 address 1.2.3.4
Wed Jun 16 21:42:38 2004 Info: Delivery start DCID 2386069 MID 200257070 to RID [0]
Wed Jun 16 21:42:38 2004 Info: ICID 282204970 close
Wed Jun 16 21:42:38 2004 Info: Message done DCID 2386069 MID 200257070 to RID [0] [('X-SBRS',
 'None')]
Wed Jun 16 21:42:38 2004 Info: MID 200257070 RID [0] Response 2.6.0
<37gva9$5uvbhe@mail.example.com> Queued mail for delivery
Wed Jun 16 21:42:43 2004 Info: DCID 2386069 close
```

Successful Message Delivery Example

```
Mon Mar 31 20:10:58 2003 Info: New SMTP DCID 5 interface 172.19.0.11 address 63.251.108.110
Mon Mar 31 20:10:58 2003 Info: Delivery start DCID 5 MID 4 to RID [0]
Mon Mar 31 20:10:58 2003 Info: Message done DCID 5 MID 4 to RID [0]
Mon Mar 31 20:11:03 2003 Info: DCID 5 close
```

Unsuccessful Message Delivery (Hard Bounce)

A message with two recipients is injected into the appliance. Upon delivery, the destination host returns a 5XX error, which indicates that the message cannot be delivered to either recipient. The appliance notifies the sender and removes the recipients from the queue.

```
Mon Mar 31 20:00:23 2003 Info: New SMTP DCID 3 interface 172.19.0.11 address 64.81.204.225
Mon Mar 31 20:00:23 2003 Info: Delivery start DCID 3 MID 4 to RID [0, 1]
Mon Mar 31 20:00:27 2003 Info: Bounced: DCID 3 MID 4 to RID 0 - 5.1.0 - Unknown address
error ('550', ['<george@yourdomain.com>... Relaying denied']) []
Mon Mar 31 20:00:27 2003 Info: Bounced: DCID 3 MID 4 to RID 1 - 5.1.0 - Unknown address
error ('550', ['<jane@yourdomain.com>... Relaying denied']) []
Mon Mar 31 20:00:32 2003 Info: DCID 3 close
```

Soft Bounce with Ultimately Successful Delivery Example

A message is injected into the appliance. On the first delivery attempt, the message soft bounces and is queued for future delivery. On the second attempt, the message is successfully delivered.

```
Mon Mar 31 20:10:58 2003 Info: New SMTP DCID 5 interface 172.19.0.11 address 63.251.108.110
Mon Mar 31 20:00:23 2003 Info: Delivery start DCID 3 MID 4 to RID [0, 1]
Mon Mar 31 20:00:23 2003 Info: Delayed: DCID 5 MID 4 to RID 0 - 4.1.0 - Unknown address
error ('466', ['Mailbox temporarily full.'])[]
Mon Mar 31 20:00:23 2003 Info: Message 4 to RID [0] pending till Mon Mar 31 20:01:23 2003
Mon Mar 31 20:01:28 2003 Info: DCID 5 close
Mon Mar 31 20:01:28 2003 Info: New SMTP DCID 16 interface PublicNet address 172.17.0.113
Mon Mar 31 20:01:28 2003 Info: Delivery start DCID 16 MID 4 to RID [0]
Mon Mar 31 20:01:28 2003 Info: Message done DCID 16 MID 4 to RID [0]
Mon Mar 31 20:01:33 2003 Info: DCID 16 close
```

Message Scanning Results (scanconfig)

When using the scanconfig command to determine behavior when a message could not be deconstructed into its component parts (when removing attachments) as with this prompt:

```
If a message could not be deconstructed into its component parts in order to remove specified
 attachments, the system should:
1. Deliver
2. Bounce
3. Drop
[3]>
```

the following is the indication in the mail logs:

With scanconfig set to deliver if message could not be decomposed.

```
Tue Aug 3 16:36:29 2004 Info: MID 256 ICID 44784 From: <test@virus.org>
Tue Aug 3 16:36:29 2004 Info: MID 256 ICID 44784 RID 0 To: <joe@example.com>
Tue Aug 3 16:36:29 2004 Info: MID 256 Message-ID '<137398.@virus.org>'
Tue Aug 3 16:36:29 2004 Info: MID 256 Subject 'Virus Scanner Test #22'
Tue Aug 3 16:36:29 2004 Info: MID 256 ready 1627 bytes from <test@virus.org>
Tue Aug 3 16:36:29 2004 Warning: MID 256, Message Scanning Problem: Continuation line seen
 before first header
Tue Aug 3 16:36:29 2004 Info: ICID 44784 close
Tue Aug 3 16:36:29 2004 Info: MID 256 antivirus positive 'EICAR-AV-Test'
Tue Aug 3 16:36:29 2004 Info: Message aborted MID 256 Dropped by antivirus
Tue Aug 3 16:36:29 2004 Info: Message finished MID 256 done
```

With scanconfig set to drop if message could not be decomposed.

```
Tue Aug 3 16:38:53 2004 Info: Start MID 257 ICID 44785
Tue Aug 3 16:38:53 2004 Info: MID 257 ICID 44785 From: test@virus.org
Tue Aug 3 16:38:53 2004 Info: MID 257 ICID 44785 RID 0 To: <joe@example.com>
Tue Aug 3 16:38:53 2004 Info: MID 257 Message-ID '<392912.@virus.org>'
Tue Aug 3 16:38:53 2004 Info: MID 25781 Subject 'Virus Scanner Test #22'
Tue Aug 3 16:38:53 2004 Info: MID 257 ready 1627 bytes from <test@virus.org>
Tue Aug 3 16:38:53 2004 Warning: MID 257, Message Scanning Problem: Continuation line seen
 before first header
Tue Aug 3 16:38:53 2004 Info: Message aborted MID 25781 Dropped by filter 'drop_zip_c'
Tue Aug 3 16:38:53 2004 Info: Message finished MID 257 done
Tue Aug 3 16:38:53 2004 Info: ICID 44785 close
```

Message with Attachment

In this example, a content filter with condition “Message Body Contains” has been configured to enable identification of attachment names:

```
Sat Apr 23 05:05:42 2011 Info: New SMTP ICID 28 interface Management (192.0.2.10)
address 224.0.0.10 reverse dns host test.com verified yes
Sat Apr 23 05:05:42 2011 Info: ICID 28 ACCEPT SG UNKNOWNLIST match sbrs[-1.0:10.0]
SBRS 0.0
Sat Apr 23 05:05:42 2011 Info: Start MID 44 ICID 28
Sat Apr 23 05:05:42 2011 Info: MID 44 ICID 28 From: <sender1@example.com>
Sat Apr 23 05:05:42 2011 Info: MID 44 ICID 28 RID 0 To: <recipient1@example.org>
Sat Apr 23 05:05:42 2011 Info: MID 44 Message-ID '<000001cba32e$f24ff2e0$d6efd8a0$@com>'
Sat Apr 23 05:05:42 2011 Info: MID 44 Subject 'Message 001'
Sat Apr 23 05:05:42 2011 Info: MID 44 ready 240129 bytes from <sender1@example.com>
Sat Apr 23 05:05:42 2011 Info: MID 44 matched all recipients for per-recipient
policy DEFAULT in the inbound table
Sat Apr 23 05:05:42 2011 Info: ICID 28 close
Sat Apr 23 05:05:42 2011 Info: MID 44 interim verdict using engine: CASE
spam negative
Sat Apr 23 05:05:42 2011 Info: MID 44 using engine: CASE spam negative
Sat Apr 23 05:05:43 2011 Info: MID 44 attachment 'Banner.gif'
Sat Apr 23 05:05:43 2011 Info: MID 44 attachment '=D1=82=D0=B5=D1=81=D1=82.rst'
Sat Apr 23 05:05:43 2011 Info: MID 44 attachment 'Test=20Attachment.docx'
Sat Apr 23 05:05:43 2011 Info: MID 44 queued for delivery
```

Note that the second of the three attachments is Unicode. On terminals that cannot display Unicode, these attachments are represented in quoted-printable format.

Generated or Rewritten Messages

Some functions, such as rewrite/redirect actions (alt-rcpt-to filters, anti-spam rcpt rewrite, bcc() actions, anti-virus redirections, and so forth), create new messages. When looking through the logs, you might need to check the results and add in additional MIDs and possibly DCIDs. Entries such as these are possible:

```
Tue Jun  1 20:02:16 2004 Info: MID 14 generated based on MID 13 by bcc filter 'nonetest'
```

or:

```
Tue Jan  6 15:03:18 2004 Info: MID 2 rewritten to 3 by antispam
Fri May 14 20:44:43 2004 Info: MID 6 rewritten to 7 by alt-rcpt-to-filter filter 'testfilt'
```

Note "Rewritten" entries can appear after lines in the log indicating use of the new MID.

Sending a Message to the Spam Quarantine

When you send a message to the quarantine, the mail logs track the movement to and from the quarantine using the RCID (RPC connection ID) to identify the RPC connection. In the following mail log, a message is tagged as spam and sent to the spam quarantine:

```
Wed Feb 14 12:11:40 2007 Info: Start MID 2317877 ICID 15726925
Wed Feb 14 12:11:40 2007 Info: MID 2317877 ICID 15726925 From: <HLD@chasehf.bfi0.com>
Wed Feb 14 12:11:40 2007 Info: MID 2317877 ICID 15726925 RID 0 To: <stevel@healthtrust.org>
Wed Feb 14 12:11:40 2007 Info: MID 2317877 Message-ID
'<W1TH05606E5811BEA0734309D4BAF0.323.14460.pimailer44.DumpShot.2@email.chase.com>'
Wed Feb 14 12:11:40 2007 Info: MID 2317877 Subject 'Envision your dream home - Now make it
 a reality'
Wed Feb 14 12:11:40 2007 Info: MID 2317877 ready 15731 bytes from <HLD@chasehf.bfi0.com>
Wed Feb 14 12:11:40 2007 Info: MID 2317877 matched all recipients for per-recipient policy
 DEFAULT in the inbound table
Wed Feb 14 12:11:41 2007 Info: MID 2317877 using engine: CASE spam suspect
Wed Feb 14 12:11:41 2007 Info: EUQ: Tagging MID 2317877 for quarantine
Wed Feb 14 12:11:41 2007 Info: MID 2317877 antivirus negative
Wed Feb 14 12:11:41 2007 Info: MID 2317877 queued for delivery
Wed Feb 14 12:11:44 2007 Info: RPC Delivery start RCID 756814 MID 2317877 to local Spam
Quarantine
Wed Feb 14 12:11:45 2007 Info: EUQ: Quarantined MID 2317877
Wed Feb 14 12:11:45 2007 Info: RPC Message done RCID 756814 MID 2317877
Wed Feb 14 12:11:45 2007 Info: Message finished MID 2317877 done
```

Using NTP Logs

The following table shows the statistics recorded in NTP logs.

Table 71: Statistics Recorded in NTP Logs

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of either a Simple Network Time Protocol (SNTP) query to the server, or an adjust: message.

Example

In this example, the NTP log shows the appliance polling the NTP host twice.

```
Thu Sep  9 07:36:39 2004 Info: sntp query host 10.1.1.23 delay 653 offset -652
Thu Sep  9 07:36:39 2004 Info: adjust: time_const: 8 offset: -652us next_poll: 4096
Thu Sep  9 08:44:59 2004 Info: sntp query host 10.1.1.23 delay 642 offset -1152
Thu Sep  9 08:44:59 2004 Info: adjust: time_const: 8 offset: -1152us next_poll: 4096
```

Using Reporting Logs

The following table shows the statistics recorded in reporting logs.

Table 72: Reporting Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken, including user authentication, and so forth.

Example

In this example, the Reporting log shows the appliance set at the information log level.

```
Wed Oct  3 13:39:53 2007 Info: Period minute using 0 (KB)
Wed Oct  3 13:39:53 2007 Info: Period month using 1328 (KB)
Wed Oct  3 13:40:02 2007 Info: Update 2 registered appliance at 2007-10-03-13-40
Wed Oct  3 13:40:53 2007 Info: Pages found in cache: 1304596 (99%). Not found: 1692
Wed Oct  3 13:40:53 2007 Info: Period hour using 36800 (KB)
Wed Oct  3 13:40:53 2007 Info: Period day using 2768 (KB)
Wed Oct  3 13:40:53 2007 Info: Period minute using 0 (KB)
Wed Oct  3 13:40:53 2007 Info: Period month using 1328 (KB)
Wed Oct  3 13:40:53 2007 Info: HELPER checkpointed in 0.00580507753533 seconds
Wed Oct  3 13:41:02 2007 Info: Update 2 registered appliance at 2007-10-03-13-41
Wed Oct  3 13:41:53 2007 Info: Pages found in cache: 1304704 (99%). Not found: 1692
Wed Oct  3 13:41:53 2007 Info: Period hour using 36800 (KB)
Wed Oct  3 13:41:53 2007 Info: Period day using 2768 (KB)
Wed Oct  3 13:41:53 2007 Info: Period minute using 0 (KB)
Wed Oct  3 13:41:53 2007 Info: Period month using 1328 (KB)
Wed Oct  3 13:42:03 2007 Info: Update 2 registered appliance at 2007-10-03-13-42
```

Using Reporting Query Logs

The following table shows the statistics recorded in reporting query logs.

Table 73: Reporting Query Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken, including user authentication, and so forth.

Example

In this example, the reporting query log shows the appliance running a daily outgoing email traffic query for the period from August 29 to October 10, 2007.

```
Tue Oct  2 11:30:02 2007 Info: Query: Closing interval handle 811804479.
```

```
Tue Oct  2 11:30:02 2007 Info: Query: Closing interval handle 811804480.
Tue Oct  2 11:30:02 2007 Info: Query: Closing query handle 302610228.
Tue Oct  2 11:30:02 2007 Info: Query: Merge query with handle 302610229 for
['MAIL_OUTGOING_TRAFFIC_SUMMARY.
DETECTED_SPAM', 'MAIL_OUTGOING_TRAFFIC_SUMMARY.DETECTED_VIRUS',
'MAIL_OUTGOING_TRAFFIC_SUMMARY.THREAT_CONTEN
T_FILTER', 'MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_CLEAN_RECIPIENTS',
'MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_RECI
PIENTS_PROCESSED'] for rollup period "day" with interval range 2007-08-29 to 2007-10-01
with key constraints
 None sorting on ['MAIL_OUTGOING_TRAFFIC_SUMMARY.DETECTED_SPAM'] returning results from 0
to 2 sort_ascendin
g=False.
Tue Oct  2 11:30:02 2007 Info: Query: Closing query handle 302610229.
Tue Oct  2 11:30:02 2007 Info: Query: Merge query with handle 302610230 for
['MAIL_OUTGOING_TRAFFIC_SUMMARY.
TOTAL_HARD_BOUNCES', 'MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_RECIPIENTS_DELIVERED',
'MAIL_OUTGOING_TRAFFIC_SUMM
ARY.TOTAL_RECIPIENTS'] for rollup period "day" with interval range 2007-08-29 to 2007-10-01
 with key constra
ints None sorting on ['MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_HARD_BOUNCES'] returning results
 from 0 to 2 sort
_ascending=False.
Tue Oct  2 11:30:02 2007 Info: Query: Closing query handle 302610230.
```

Using Safelist/Blocklist Logs

The following table shows the statistics recorded in safelist/blocklist logs.

Table 74: Safelist/Blocklist Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken, including user authentication, and so forth.

Example

In this example, the safelist/blocklist log shows the appliance creating database snapshots every two hours. It also shows when senders were added to the database.

```
Fri Sep 28 14:22:33 2007 Info: Begin Logfile Fri Sep 28 14:22:33 2007 Info: Version: 6.0.0-425
 SN: XXXXXXXXXXXX-XXX Fri Sep 28 14:22:33 2007 Info: Time offset from UTC: 10800 seconds
Fri Sep 28 14:22:33 2007 Info: System is coming up.
Fri Sep 28 14:22:33 2007 Info: SLBL: The database snapshot has been created.
Fri Sep 28 16:22:34 2007 Info: SLBL: The database snapshot has been created.
Fri Sep 28 18:22:34 2007 Info: SLBL: The database snapshot has been created.
Fri Sep 28 20:22:34 2007 Info: SLBL: The database snapshot has been created.
Fri Sep 28 22:22:35 2007 Info: SLBL: The database snapshot has been created.
.........................
Mon Oct  1 14:16:09 2007 Info: SLBL: The database snapshot has been created.
Mon Oct  1 14:37:39 2007 Info: SLBL: The database snapshot has been created.
Mon Oct  1 15:31:37 2007 Warning: SLBL: Adding senders to the database failed.
Mon Oct  1 15:32:31 2007 Warning: SLBL: Adding senders to the database failed.
Mon Oct  1 16:37:40 2007 Info: SLBL: The database snapshot has been created.
```

Using SMA Logs

The following table shows the statistics recorded in SMA logs.

Table 75: SMA Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken, including user authentication, and so forth.

Example

In this example, the SMA log shows the centralized tracking service downloading tracking files from an Email Security appliance, and it shows the centralized reporting service downloading reporting files from an Email Security appliance.

```
Wed Oct  3 13:26:39 2007 Info: TRANSFER: Plugin TRACKINGPLUGIN downloading from 172.29.0.17
 - /export/tracki
ng/tracking.@20071003T202244Z_20071003T202544Z.s
Wed Oct  3 13:28:11 2007 Info: TRANSFER: Plugin TRACKINGPLUGIN downloading from 172.29.0.15
 - /export/tracki
ng/tracking.@20071003T202443Z_20071003T202743Z.s
Wed Oct  3 13:28:46 2007 Info: TRANSFER: Plugin TRACKINGPLUGIN downloading from 172.29.0.17
 - /export/tracki
ng/tracking.@20071003T202544Z_20071003T202844Z.s
Wed Oct  3 13:31:27 2007 Info: TRANSFER: Plugin TRACKINGPLUGIN downloading from 172.29.0.15
 - /export/tracki
ng/tracking.@20071003T202743Z_20071003T203043Z.s
Wed Oct  3 13:31:28 2007 Info: TRANSFER: Plugin REPORTINGPLUGIN downloading from 172.29.0.15
 - /reporting/ou
tgoing_queue/rpx.2007-10-03-20-15Z.000F1F6ECA7C-2RWDB51.v1.tgz
Wed Oct  3 13:31:53 2007 Info: TRANSFER: Plugin TRACKINGPLUGIN downloading from 172.29.0.17
 - /export/tracki
ng/tracking.@20071003T202844Z_20071003T203144Z.s
Wed Oct  3 13:32:31 2007 Info: TRANSFER: Plugin REPORTINGPLUGIN downloading from 172.29.0.17
 - /reporting/ou
tgoing_queue/rpx.2007-10-03-20-15Z.0019B9B316E4-JZ41PC1.v1.tgz
Wed Oct  3 13:34:40 2007 Info: TRANSFER: Plugin TRACKINGPLUGIN downloading from 172.29.0.15
 - /export/tracki
ng/tracking.@20071003T203043Z_20071003T203343Z.s
```

Using Status Logs

Status logs record system statistics found in the CLI status commands, including status , status detail , and dnsstatus . The period of recording is set using the setup subcommand in logconfig . Each counter or rate reported in status logs is the value since the last time the counter was reset.

Table 76: Status Log Statistics

Statistic	Description
CPULd	CPU utilization.

Statistic	Description
DskIO	Disk I/O utilization.
RAMUtil	RAM utilization.
QKUsd	Queue kilobytes used.
QKFre	Queue kilobytes free.
CrtMID	Message ID (MID).
CrtICID	Injection connection ID (ICID).
CRTDCID	Delivery connection ID (DCID).
InjMsg	Injected messages.
InjRcp	Injected recipients.
GenBncRcp	Generated bounce recipients.
RejRcp	Rejected recipients.
DrpMsg	Dropped messages.
SftBncEvnt	Soft bounced events.
CmpRcp	Completed recipients.
HrdBncRcp	Hard bounced recipients.
DnsHrdBnc	DNS hard bounces.
5XXHrdBnc	5XX hard bounces.
FltrHrdBnc	Filter hard bounces.
ExpHrdBnc	Expired hard bounces.
OtrHrdBnc	Other hard bounces.
DlvRcp	Delivered recipients.
DelRcp	Deleted recipients.
GlbUnsbHt	Global unsubscribe hits.
ActvRcp	Active recipients.

Statistic	Description
UnatmptRcp	Unattempted recipients.
AtmptRcp	Attempted recipients.
CrtCncIn	Current inbound connections.
CrtCncOut	Current outbound connections.
DnsReq	DNS requests.
NetReq	Network requests.
CchHit	Cache hits.
CchMis	Cache misses.
CchEct	Cache exceptions.
CchExp	Cache expired.
CPUTTm	Total CPU time used by the application.
CPUETm	Elapsed time since the application started.
MaxIO	Maximum disk I/O operations per second for the mail process.
RamUsd	Allocated memory in bytes.
SwIn	Memory swapped in.
SwOut	Memory swapped out.
SwPgIn	Memory paged in.
SwPgOut	Memory paged out.
MMLen	Total number of messages in the system.
DstInMem	Number of destination objects in memory.
ResCon	Resource conservation tarpit value. Acceptance of incoming mail is delayed by this number of seconds due to heavy system load.
WorkQ	Number of messages currently in the work queue.

Statistic	Description
QuarMsgs	Number of individual messages in the system quarantine (messages present in multiple quarantines are counted only once).
QuarQKUsd	Kilobytes used by system quarantine messages.
LogUsd	Percent of log partition used.
CASELd	Percent CPU used by CASE scanning.
TotalLd	Total CPU consumption.
LogAvail	Amount of disk space available for log files.
EuQ	Number of messages in the spam quarantine.
EuqRls	Number of messages in the spam quarantine release queue.

Example

```
Fri Feb 24 15:14:39 2006 Info: Status: CPULd 0 DskIO 0 RAMUtil 2 QKUsd 0 QKFre 8388608
CrtMID 19036 CrtICID 35284 CrtDCID 4861 InjMsg 13889 InjRcp 14230 GenBncRcp 12 RejRcp 6318
 DrpMsg 7437 SftBncEvnt 1816 CmpRcp 6813 HrdBncRcp 18 DnsHrdBnc 2 5XXHrdBnc 15 FltrHrdBnc
0 ExpHrdBnc 1 OtrHrdBnc 0 DlvRcp 6793 DelRcp 2 GlbUnsbHt 0 ActvRcp 0 UnatmptRcp 0 AtmptRcp
 0 CrtCncIn 0 CrtCncOut 0 DnsReq 143736 NetReq 224227 CchHit 469058 CchMis 504791 CchEct
15395 CchExp 55085 CPUTTm 228 CPUETm 181380 MaxIO 350 RAMUsd 21528056 MMLen 0 DstInMem 4
ResCon 0 WorkQ 0 QuarMsgs 0 QuarQKUsd 0 LogUsd 3 AVLd 0 BMLd 0 CASELd 3 TotalLd 3 LogAvail
 17G EuQ 0 EuqRls 0
```

Using System Logs

The following table shows the statistics recorded in system logs.

Table 77: System Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The logged event.

Example

In this example, the system log shows some commit entries, including the name of the user issuing the commit and the comment entered.

```
Wed Sep  8 18:02:45 2004 Info: Version: 6.0.0-206 SN: XXXXXXXXXXXX-XXX
Wed Sep  8 18:02:45 2004 Info: Time offset from UTC: 0 seconds
```

```
Wed Sep  8 18:02:45 2004 Info: System is coming up
Wed Sep  8 18:02:49 2004 Info: bootstrapping DNS cache
Wed Sep  8 18:02:49 2004 Info: DNS cache bootstrapped
Wed Sep  8 18:13:30 2004 Info: PID 608: User admin commit changes: SSW:Password
Wed Sep  8 18:17:23 2004 Info: PID 608: User admin commit changes: Completed Web::SSW
Thu Sep  9 08:49:27 2004 Info: Time offset from UTC: -25200 seconds
Thu Sep  9 08:49:27 2004 Info: PID 1237: User admin commit changes: Added a second CLI log
 for examples
Thu Sep  9 08:51:53 2004 Info: PID 1237: User admin commit changes: Removed example CLI
log.
```

Understanding Tracking Logs

Tracking logs record information about the email operations of AsyncOS. The log messages are a subset of the messages recorded in the mail logs.

The tracking logs are used by the message tracking component to build the message tracking database. Because the log files are consumed in the process of building the database, the tracking logs are transient. The information in tracking logs is not designed to be read or analyzed by humans.

Tracking logs are recorded and transferred in a binary format for resource efficiency. The information is laid out in a logical manner and is human-readable after conversion using a utility provided by Cisco. The conversion tools are located at the following URL: http://tinyurl.com/3c5l8r .

Log Subscriptions

Configuring Log Subscriptions

Log subscriptions create the individual log files that are stored on a Cisco Content Security appliance or remotely. A log subscription is either pushed (delivered to another computer) or polled (retrieved from the appliance). Generally, log subscriptions have the following attributes:

Table 78: Log File Attributes

Attribute	Description
Log Type	Defines the type of information recorded and the format of the log subscription. For more information, see Summary of Log Types, on page 350.
Name	Descriptive name of log subscription that you provide for your future reference.
Log Filename	Physical name of the file when it is written to disk. If the system includes multiple content security appliances, use a unique log filename to identify the appliance that generated the log file.
Rollover by File Size	Maximum size that the file can reach before it rolls over.
Rollover by Time	When to roll over log files, based on time. See options at Log Rollover and Transfer Schedule, on page 349.

Attribute	Description
Log Level	Level of detail for each log subscription.
Retrieval Method	Method used to transfer the log files from the appliance.

Use the Management Appliance > System Administration > Log Subscriptions page (or the logconfig command in the CLI) to configure a log subscription. You are prompted for the log type, as shown in Summary of Log Types, on page 350. For most log types, you are also asked to select a *log level* for the log subscription.

Note Configuration history logs only: If you anticipate loading configurations from the configuration history logs, be aware that you cannot load configurations containing masked passwords. On the Management Appliance > System Administration > Log Subscriptions page, select Yes when prompted whether you want to include passwords in the log. If you are using the logconfig command in the CLI, type y when prompted.

Setting the Log Level

Log levels determine the amount of information delivered in a log. Logs can have one of five levels of detail. A detailed log-level setting creates larger log files and has a greater impact on system performance than an abbreviated log-level setting. A detailed log-level setting includes all the messages contained in the abbreviated log-level settings, plus additional messages. As the level of detail increases, system performance decreases.

Note You can specify different logging levels for each log type.

Table 79: Log Levels

Log Level	Description
Critical	Only errors are logged. This is the most abbreviated log-level setting. At this log level, you cannot monitor performance and important appliance activities; however, the log files do not reach maximum size as quickly as they do at a detailed log level. This log level is analogous to the syslog level Alert.
Warning	All system errors and warnings are logged. At this log level, you cannot monitor performance and important appliance activities. The log files reach maximum size more quickly than they do at the Critical log level. This log level is analogous to the syslog level Warning.

Log Level	Description
Information	Second-by-second operations of the system are logged. For example, connections opened and delivery attempts are logged. The Information level is the recommended setting for logs. This log level is analogous to the syslog level Info.
Debug	More detailed information is logged than at the Information log level. Use the Debug log level when you are troubleshooting an error. Use this setting temporarily, and then return to the default level. This log level is analogous to the syslog level Debug.
Trace	All available information is logged. The Trace log level is recommended only for developers. Using this level causes a serious degradation of system performance and is not recommended. This log level is analogous to the syslog level Debug.

Creating a Log Subscription in the GUI

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 On the **Management Appliance > System Administration > Log Subscriptions** page, click **Add Log Subscription**.

Step 3 Select a log type and enter the log name (for the log directory), as well as the name for the log file itself.

Step 4 If applicable, specify the maximum file size.

Step 5 If applicable, specify days, times of day, or time intervals to roll over the logs. For more information, see Log Rollover and Transfer Schedule, on page 349.

Step 6 If applicable, specify the log level.

Step 7 (Configuration history logs only) Select whether to include passwords in the log.

Note You cannot load configurations containing masked passwords. If you anticipate loading configurations from the configuration history logs, select Yes to include passwords in the log.

Step 8 Configure the log retrieval method.

Step 9 Submit and commit your changes.

Editing Log Subscriptions

Step 1 Click the name of the log in the Log Name column on the Log Subscriptions page.

Step 2 Update the log subscription.

Step 3 Submit and commit your changes.

Configuring Global Settings for Logging

The system periodically records system metrics within text mail logs and status logs. Use the **Edit Settings** button in the Global Settings section of the Log Subscriptions page (or the logconfig -> setup command in the CLI) to configure:

- The amount of time, in seconds, that the system waits between recording metrics
- Whether to record the Message ID headers
- Whether to record the remote response status code
- Whether to record the subject header of the original message
- The headers that should be logged for each message

All Cisco Content Security appliance logs optionally include the following three items:

- Message-ID: When this option is configured, every message will have its Message ID header logged, if it is available. This Message ID may have come from the received message or may have been generated by AsyncOS. For example:

```
Tue Apr 6 14:38:34 2004 Info: MID 1 Message-ID Message-ID-Content
```

- Remote Response: When this option is configured, every message will have its remote response status code logged, if it is available. For example:

```
Tue Apr 6 14:38:34 2004 Info: MID 1 RID [0] Response 'queued as 9C8B425DA7'
```

The remote response string is the human-readable text received after the response to the DATA command during the delivery SMTP conversation. In this example, the remote response after the connection host issued the data command is “queued as 9C8B425DA7.”

```
[...]
250 ok hostname
250 Ok: queued as 9C8B425DA7
```

White space, punctuation, and, in the case of the 250 response, the OK characters are stripped from the beginning of the string. Only white space is stripped from the end of the string. For example, Cisco Content Security appliances, by default, respond to the DATA command with this string: 250 Ok: Message MID accepted . So, the entry “Message MID accepted” would be logged if the remote host were another Cisco Content Security appliance.

- Original Subject Header: When this option is enabled, the original subject header of each message is included in the log.

```
Tue May 31 09:20:27 2005 Info: Start MID 2 ICID 2
Tue May 31 09:20:27 2005 Info: MID 2 ICID 2 From: <mary@example.com>
Tue May 31 09:20:27 2005 Info: MID 2 ICID 2 RID 0 To: <joe@example.com>
Tue May 31 09:20:27 2005 Info: MID 2 Message-ID '<44e4n$2@example.com>'
Tue May 31 09:20:27 2005 Info: MID 2 Subject 'Monthly Reports Due'
```

Logging Message Headers

In some cases, it is necessary to record the presence and contents of a message's headers as they pass through the system. You specify the headers to record on the Log Subscriptions Global Settings page (or via the logconfig -> logheaders subcommand in the CLI). The appliance records the specified message headers in the text mail logs and the tracking logs. If the header is present, the system records the name of the header and the value. If a header is not present, nothing is recorded in the logs.

Note The system evaluates all headers that are present on a message, at any time during the processing of the message for recording, regardless of the headers specified for logging.

Note The RFC for the SMTP protocol is located at http://www.faqs.org/rfcs/rfc2821.html and defines user-defined headers.

Note If you have configured headers to log via the logheaders command, the header information appears after the delivery information:

Table 80: Log Headers

Header name	Name of the header
Value	Contents of the logged header

For example, specifying "date, x-subject" as headers to be logged causes the following line to appear in the mail log:

```
Tue May 31 10:14:12 2005 Info: Message done DCID 0 MID 3 to RID [0] [('date', 'Tue, 31 May
 2005 10:13:18 -0700'), ('x-subject', 'Logging this header')]
```

Configuring Global Settings for Logging by Using the GUI

Step 1 Click the **Edit Settings** button in the Global Settings section of the Log Subscriptions page.

Step 2 Specify the system metrics frequency, whether to include Message ID headers in mail logs, whether to include the remote response, and whether to include the original subject header of each message.
For information about these settings, see Configuring Global Settings for Logging, on page 374.

Step 3 Enter any other headers you want to include in the logs. Separate each entry with a comma.

Step 4 Submit and commit your changes.

Rolling Over Log Subscriptions

When AsyncOS rolls over a log file, it:

- Creates a new log file with the timestamp of the rollover and designates the file as current with the letter "**c**" extension
- Renames the current log file to have a letter "**s**" extension signifying saved
- Transfers the newly saved log file to a remote host (if push-based)
- Transfers any previously unsuccessful log files from the same subscription (if push-based)
- Deletes the oldest file in the log subscription if the total number of files to keep on hand has been exceeded (if poll-based)

What To Do Next

Rolling Over Logs in Log Subscriptions

See Log Rollover and Transfer Schedule, on page 349.

Rolling Over Logs Immediately Using the GUI

Step 1 On the Log Subscriptions page, select the check box to the right of the logs you want to roll over.

Step 2 Optionally, select all logs for rollover by selecting the **All** check box.

Step 3 Click the **Rollover Now** button.

What to Do Next

- Rolling Over Logs in Log Subscriptions , on page 376
- Rolling Over Logs Immediately via the CLI, on page 376

Rolling Over Logs Immediately via the CLI

Use the rollovernow command to roll over all log files at once or select a specific log file from a list.

Viewing the Most Recent Log Entries in the GUI

You can view a log file via the GUI by clicking the log subscription in the Log Files column of the table on the Log Subscriptions page. When you click the link to the log subscription, you are prompted to enter your password. A listing of log files for that subscription then appears. You can click one of the log files to view it in your browser or to save it to disk. You must have the FTP service enabled on the Management interface to view logs in the GUI.

Viewing the Most Recent Entries in Logs (tail Command)

AsyncOS supports a tail command, which shows the latest entries of configured logs on the appliance. Issue the tail command and select the number of a currently configured log to view it. Press Ctrl-C to exit from the tail command.

Note You cannot view configuration history logs by using the tail command. You must use FTP or SCP.

Example

In the following example, the tail command is used to view the system log. The tail command also accepts the name of a log to view as a parameter, for example, tail system_logs

```
Welcome to the M600 Messaging Gateway(tm) Appliance
example.srv> tail
Currently configured logs:
1. "cli_logs" Type: "CLI Audit Logs" Retrieval: FTP Poll
2. "euq_logs" Type: " Spam Quarantine Logs" Retrieval: FTP Poll
3. "euqgui_logs" Type: "Spam Quarantine GUI Logs" Retrieval: FTP Poll
4. "gui_logs" Type: "HTTP Logs" Retrieval: FTP Poll
5. "mail_logs" Type: "Text Mail Logs" Retrieval: FTP Poll
6. "reportd_logs" Type: "Reporting Logs" Retrieval: FTP Poll
7. "reportqueryd_logs" Type: "Reporting Query Logs" Retrieval: FTP Poll
8. "slbld_logs" Type: "Safe/Block Lists Logs" Retrieval: FTP Poll
9. "smad_logs" Type: "SMA Logs" Retrieval: FTP Poll
10. "system_logs" Type: "System Logs" Retrieval: FTP Poll
11. "trackerd_logs" Type: "Tracking Logs" Retrieval: FTP Poll
Enter the number of the log you wish to tail.
[]> 10
Press Ctrl-C to stop.
Thu Sep 27 00:18:56 2007 Info: Begin Logfile
Thu Sep 27 00:18:56 2007 Info: Version: 6.0.0-422 SN: 001143583D73-FT9GP61
Thu Sep 27 00:18:56 2007 Info: Time offset from UTC: 0 seconds
Thu Sep 27 00:18:47 2007 Info: System is coming up.
Thu Sep 27 00:23:05 2007 Warning: DNS query network error '[Errno 64] Host is down' to
'172.16.0.3' looking up 'downloads.cisco.com'
Fri Sep 28 22:20:08 2007 Info: PID 688: User admin commit changes:
Fri Sep 28 23:06:15 2007 Info: PID 688: User admin commit changes:
^Cexample.srv>
.
```

Configuring Host Keys

Use the logconfig -> hostkeyconfig subcommand to manage host keys for use with SSH when pushing logs to other servers from the Cisco Content Security appliance. SSH servers must have a pair of host keys, one private and one public. The private host key resides on the SSH server and cannot be read by remote machines. The public host key is distributed to any client machine that needs to interact with the SSH server.

Note To manage user keys, see “Managing Secure Shell (SSH) Keys” in the user guide or online help for your Email Security appliance.

The hostkeyconfig subcommand performs the following functions:

Table 81: Managing Host Keys - List of Subcommands

Command	Description
New	Add a new key.
Edit	Modify an existing key.
Delete	Delete an existing key.
Scan	Automatically download a host key.
Print	Display a key.
Host	Display system host keys. This is the value to place in the remote system's “known_hosts” file.
Fingerprint	Display system host key fingerprints.
User	Display the public key of the system account that pushes the logs to the remote machine. This is the same key that appears when setting up an SCP push subscription. This is the value to place in the remote system's “authorized_keys” file.

Example

In the following example, the commands scan for host keys and add them for the host:

```
mail3.example.com> logconfig
Currently configured logs:
[ list of logs
 ]
Choose the operation you want to perform:
- NEW - Create a new log.
- EDIT - Modify a log subscription.
- DELETE - Remove a log subscription.
- SETUP - General settings.
- LOGHEADERS - Configure headers to log.
- HOSTKEYCONFIG - Configure SSH host keys.
[]> hostkeyconfig
Currently installed host keys:
1. mail3.example.com ssh-dss [ key displayed ]
Choose the operation you want to perform:
- NEW - Add a new key.
- EDIT - Modify a key.
- DELETE - Remove a key.
- SCAN - Automatically download a host key.
- PRINT - Display a key.
- HOST - Display system host keys.
- FINGERPRINT - Display system host key fingerprints.
- USER - Display system user keys.
```

```
[]> scan
Please enter the host or IP address to lookup.
[]> mail3.example.com
Choose the ssh protocol type:
1. SSH2:rsa
2. SSH2:dsa
3. All
[3]>
SSH2:dsa
mail3.example.com ssh-dss
[ key displayed
 ]
SSH2:rsa
mail3.example.com ssh-rsa
[ key displayed
 ]
Add the preceding host key(s) for mail3.example.com? [Y]>
Currently installed host keys:
1. mail3.example.com ssh-dss [ key displayed
 ]
2. mail3.example.com ssh-rsa [ key displayed
 ]
3. mail3.example.com 1024 35 [ key displayed
 ]
Choose the operation you want to perform:
- NEW - Add a new key.
- EDIT - Modify a key.
- DELETE - Remove a key.
- SCAN - Automatically download a host key.
- PRINT - Display a key.
- HOST - Display system host keys.
- FINGERPRINT - Display system host key fingerprints.
- USER - Display system user keys.
[]>
Currently configured logs:
[ list of configured logs
 ]
Choose the operation you want to perform:
- NEW - Create a new log.
- EDIT - Modify a log subscription.
- DELETE - Remove a log subscription.
- SETUP - General settings.
- LOGHEADERS - Configure headers to log.
- HOSTKEYCONFIG - Configure SSH host keys.
[]>
mail3.example.com> commit
```

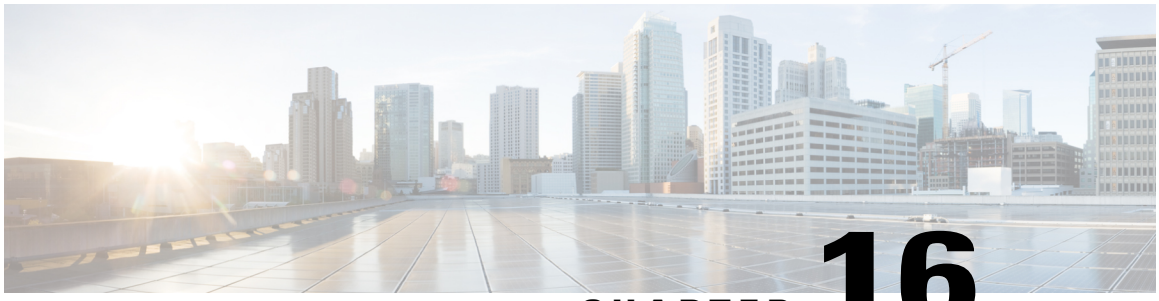

CHAPTER 16

Troubleshooting

This chapter contains the following sections:

Collecting System Information

You can get information about your appliance and its status, including your serial number. Refer Monitoring System Status, on page 223

Troubleshooting Hardware Issues

Lights on the front and/or rear panels of your hardware appliance indicate health and status of your appliance. For descriptions of these indicators, see the hardware guides such as the *Cisco x90 Series Content Security Appliances Installation and Maintenance Guide* available from the location specified in.

Specifications for your appliance, such as temperature ranges, are also available in these documents.

Note If you need to cycle power to your x80 or x90 appliance, wait at least 20 minutes for the appliance to come up (all LEDs are green) before pushing the power button.

Troubleshooting Feature Setup Issues

If you are experiencing difficulty configuring a feature successfully, see the summaries of the tasks you must complete for each feature. These include links to specific information for each.

- Setting Up Centralized Email Reporting, on page 36
- Setting Up Centralized Message Tracking , on page 160
- Setting Up the Centralized Spam Quarantine , on page 174
- Centralized Policy, Virus, and Outbreak Quarantines, on page 197

General Troubleshooting Resources

General troubleshooting resources include:

- Recent alerts. See Viewing Recent Alerts , on page 314.
- Log files. SeeLogging, on page 347
- The Release Notes, including the Documentation Updates section. SeeDocumentation , on page 405.
- The Cisco Bug Search Tool (instructions for access are in the Release Notes)
- Knowledge Base Articles (TechNotes) , on page 407
- TheCisco Support Community, on page 407

Troubleshooting Issues with Specific Functionality

See alsoTroubleshooting Feature Setup Issues , on page 382.

Email Security-Related Issues

- Troubleshooting All Reports , on page 33
- Troubleshooting Message Tracking, on page 172
- Troubleshooting Spam Quarantine Features , on page 196
- Troubleshooting Centralized Policy Quarantines , on page 222
- Feature-related issues may also result from settings on your Email Security appliances. See the release notes and online help or user guide for your release at the location specified inDocumentation , on page 405.

General Issues

- If you are unable to load a configuration file, make sure your disk space quotas are larger than the current size of each function in the table on the **Management Appliance > System Administration > Disk Management** page.

- If you have recently upgraded and the online help appears to be outdated or you cannot find the information about a new feature, clear your browser cache and then reopen the browser window.
- Unexpected behavior can occur when configuring settings using the web interface if you are using multiple browser windows or tabs simultaneously.
- SeeResponding to Alerts , on page 383.
- SeeTroubleshooting Administrative User Access , on page 278.

Responding to Alerts

Alert: Battery Relearn Timed Out (RAID Event) on 380 or 680 Hardware

Problem: You receive an alert with subject “Battery Relearn Timed Out” for 380 or 680 hardware.

Solution: This alert may or may not indicate a problem. The battery relearn timeout, in itself, does not mean there is any problem with the RAID controller. The controller can recover in the subsequent relearn. Please monitor your email for any other RAID alerts for the next 48 hours, to ensure that this is not the side-effect of any other problem. If you do not see any other RAID-related alerts from the system, then you can safely ignore this alert.

Additional Alert Descriptions

For descriptions of additional alerts, see

- Hardware Alert Descriptions , on page 315
- System Alert Descriptions , on page 316

What to do next

- Managing Alerts, on page 312

Working with Technical Support

Opening or Updating a Support Case from the Appliance

You can use this method to contact Cisco TAC or your own support services.

Before You Begin

If you wish to contact Cisco TAC:

- If your issue is urgent, do not use this method. Instead, contact support using one of the other methods listed in Customer Support , on page 407.
- Consider other options for getting help:

- When you open a support case using this procedure, the appliance configuration file is sent to Cisco Customer Support. If you do not want to send the appliance configuration, you can contact Customer Support using a different method.
- The appliance must be connected to the internet and able to send email.
- If you are sending information about an existing case, make sure you have the case number.

Step 1 Log in to the appliance.

Step 2 Choose **Help and Support > Contact Technical Support**.

Step 3 Determine the recipients of the support request:

To send the request to Cisco TAC	Select the **Cisco Technical Support** check box.
To send the request only to your internal support desk	• Deselect the **Cisco Technical Support** check box. • Enter the email address of your support desk.
(Optional) To include other recipients	Enter email addresses.

Step 4 Complete the form.

Step 5 Click **Send**.

Getting Support for Virtual Appliances

If you file a support case for a Cisco content security virtual appliance, you must provide your Virtual License Number (VLN), your contract number, and your Product Identifier code (PID).

You can identify your PID based on the software licenses running on your virtual appliance, by referencing your purchase order, or from the following table:

Functionality	PID	Description
All centralized email security functionality	SMA-EMGT-LIC=	

Enabling Remote Access for Cisco Technical Support Personnel

Only Cisco Customer Assistance can access your appliance using these methods.

Enabling Remote Access to Appliances With an Internet Connection

Support accesses the appliance through an SSH tunnel that this procedure creates between the appliance and the upgrades.ironport.com server.

Before You Begin

Identify a port that can be reached from the internet. The default is port25 which will work in most environment . Connections over this port are allowed in most firewall configurations.

Step 1 Log in to the appliance.

Step 2 From the top right side of the GUI window, choose **Help and Support > Remote Access**.

Step 3 Click **Enable**.

Step 4 Enter information.

Step 5 Click **Submit**.

What to Do Next

When remote access for support personnel is no longer required, seeDisabling a Tech Support Tunnel , on page 386.

Enabling Remote Access to Appliances Without a Direct Internet Connection

For appliances without a direct internet connection, access is made through a second appliance that is connected to the internet.

Before You Begin

- The appliance must be able to connect on port 22 to a second appliance that is connected to the internet.
- On the appliance with the internet connection, follow the procedure inEnabling Remote Access to Appliances With an Internet Connection , on page 385 to create a support tunnel to that appliance.

Step 1 From the command-line interface of the appliance requiring support, enter the `techsupport` command.

Step 2 Enter `sshaccess` .

Step 3 Follow the prompts.

What to Do Next

When remote access for support personnel is no longer required, see the following:

- Disabling Remote Access , on page 386
- Disabling a Tech Support Tunnel , on page 386

Disabling a Tech Support Tunnel

An enabled techsupport tunnel remains connected to upgrades.ironport.com for 7 days. After that time, established connections will not be disconnected but will be unable to re-attach to the tunnel once disconnected.

Step 1 Log in to the appliance.

Step 2 From the top right side of the GUI window, choose **Help and Support > Remote Access**.

Step 3 Click **Disable**.

Disabling Remote Access

A remote access account that you create using the techsupport command remains active until you deactivate it.

Step 1 From the command-line interface, enter the `techsupport` command.

Step 2 Enter `sshaccess` .

Step 3 Enter `disable` .

Checking the Status of the Support Connection

Step 1 From the command-line interface, enter the `techsupport` command.

Step 2 Enter status .

Running a Packet Capture

Packet Capture allows support personnel to see the TCP/IP data and other packets going into and out of the appliance. This allows Support to debug the network setup and to discover what network traffic is reaching the appliance or leaving the appliance.

Step 1 Choose **Help and Support > Packet Capture**.

Step 2 Specify packet capture settings:

a) In the **Packet Capture Settings** section, click **Edit Settings**.

b) (Optional) Enter duration, limits, and filters for the packet capture.
Your Support representative may give you guidance on these settings.

If you enter a capture duration without specifying a unit of time, AsyncOS uses seconds by default.

In the Filters section:

- Custom filters can use any syntax supported by the Unix `tcpdump` command, such as host 10.10.10.10 && port 80 .
- The client IP is the IP address of the machine connecting to the appliance, such as a mail client sending messages through the Email Security appliance.
- The server IP is the IP address of the machine to which the appliance is connecting, such as an Exchange server to which the appliance is delivering messages.

You can use the client and server IP addresses to track traffic between a specific client and a specific server, with the Email Security appliance in the middle.

c) Click **Submit**.

Step 3 Click **Start Capture**.

- Only one capture may be running at a time.
- When a packet capture is running, the Packet Capture page shows the status of the capture in progress by showing the current statistics, such as file size and time elapsed.
- The GUI only displays packet captures started in the GUI, not from the CLI. Similarly, the CLI only displays the status of a current packet capture run started in the CLI.
- The packet capture file is split into ten parts. If the file reaches the maximum size limit before the packet capture ends, the oldest part of the file is deleted (the data is discarded) and a new part starts with the current packet capture data. Only 1/10 of the packet capture file is discarded at a time.
- A running capture started in the GUI is preserved between sessions. (A running capture started in the CLI stops when the session ends.)

Step 4 Allow the capture to run for the specified duration, or, if you have let the capture run indefinitely, manually stop the capture by clicking **Stop Capture**.

Step 5 Access the packet capture file:

- Click the file in the **Manage Packet Capture Files** list and click **Download File**.
- Use FTP or SCP to access the file in the captures subdirectory on the appliance.

What to Do Next

Make the file available to Support:

- If you allow remote access to your appliance, technicians can access the packet capture files using FTP or SCP. SeeEnabling Remote Access for Cisco Technical Support Personnel , on page 384.
- Email the file to Support.

Remotely Resetting Appliance Power

If the appliance requires a hard reset, you can reboot the appliance chassis remotely using a third-party Intelligent Platform Management Interface (IPMI) tool.

Restrictions

- Remote power cycling is available only on certain hardware.

  For specifics, seeEnabling Remote Power Cycling , on page 287.
- If you want be able to use this feature, you must enable it in advance.

  For details, seeEnabling Remote Power Cycling , on page 287.
- Only the following IPMI commands are supported:

  `status, on, off, cycle, reset, diag, soft`

  Issuing unsupported commands will produce an “insufficient privileges” error.

Before You Begin

- Obtain and set up a utility that can manage devices using IPMI version 2.0.
- Understand how to use the supported IPMI commands. See the documentation for your IPMI tool.

Step 1 Use IPMI to issue a supported power-cycling command to the IP address assigned to the Remote Power Cycle port, which you configured earlier, along with the required credentials.
For example, from a UNIX-type machine with IPMI support, you might issue the command:

```
ipmitool -I lan -H 192.0.2.1 -U remoteresetuser -P password chassis power reset
```

where `192.0.2.1` is the IP address assigned to the Remote Power Cycle port and remoteresetuser and password are the credentials that you entered while enabling this feature.

Step 2 Wait at least eleven minutes for the appliance to reboot.

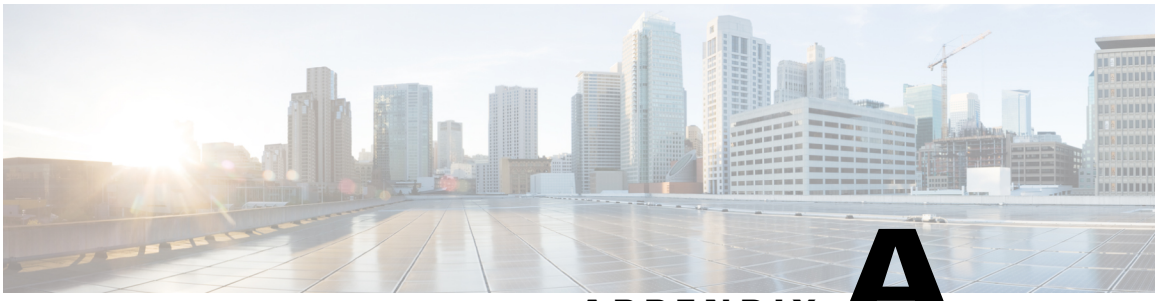

APPENDIX A

IP Interfaces and Accessing the Appliance

This chapter contains the following sections:

IP Interfaces and Accessing the Appliance

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

Table 82: Services Enabled by Default on IP Interfaces

		Enabled by default?	
Service	**Default Port**	**Management Interface**	**New IP Interfaces You Create**
FTP	21	No	No
Telnet	23	Yes	No
SSH	22	Yes	No
HTTP	80	Yes	No
HTTPS	443	Yes	No

IP Interfaces

An IP interface contains the network configuration data needed for an individual connection to the network. You can configure multiple IP interfaces to a physical Ethernet interface. You can also configure access to

the spam quarantine via an IP interface. For email delivery and Virtual Gateways, each IP interface acts as one Virtual Gateway address with a specific IP address and hostname. You can also "join" interfaces into distinct groups (via the CLI), and the system will cycle through these groups when delivering email. Joining or grouping Virtual Gateways is useful for load-balancing large email campaigns across several interfaces. You can also create VLANs, and configure them just as you would any other interface (via the CLI). For more information, see the "Advanced Networking" chapter in the user guide or online help for your Email Security appliance.

Configuring IP Interfaces

The Management Appliance > Network > IP Interfaces page (and interface config command) enables you to add, edit, or delete IP interfaces.

Note You cannot change the name or Ethernet port associated with the Management interface on the Security Management appliance. Further, the Security Management appliance does not support all of the features discussed below (Virtual Gateways, for example).

The following information is required when you configure an IP interface:

Table 83: IP Interface Components

Name	**The nickname of the interface.**
IP address	IP addresses within the same subnet cannot be configured on separate physical Ethernet interfaces.
Netmask (or subnetmask)	You can enter the netmask in standard dotted octet form (for example, 255.255.255.0) or hexadecimal form (for example, 0xffffff00). The default netmask is 255.255.255.0, a common class C value.
Broadcast address	AsyncOS automatically calculates the default broadcast address from the IP address and the netmask.
Hostname	The hostname that is related to the interface. This hostname is used to identify the server during the SMTP conversation. You are responsible for entering a valid hostname associated with each IP address. The software does not check that DNS correctly resolves the hostname to the matching IP address, or that reverse DNS resolves to the given hostname.
Allowed services	FTP, SSH, Telnet, spam quarantine, HTTP, and HTTPS can be enabled or disabled on the interface. You can configure the port for each service. You can also specify the HTTP/HTTPS, port, and URL for the spam quarantine.

Note If you have completed the System Setup Wizard as described in Setup, Installation, and Basic Configuration, *on page 9* and committed the changes, the Management interface should already be configured on the appliance.

Creating IP Interfaces Using the GUI

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance > Network > IP Interfaces**.

Step 3 Click **Add IP Interface**.

Step 4 Enter a name for the interface.

Step 5 Select an Ethernet port and enter an IP address.

Step 6 Enter the netmask for the IP address.

Step 7 Enter a hostname for the interface.

Step 8 Select the check box next to each service you want to enable on this IP interface. Change the corresponding port if necessary.

Step 9 Select whether to enable redirecting HTTP to HTTPS for appliance management on the interface.

Step 10 If you are using the spam quarantine, you can select HTTP or HTTPS or both and specify the port numbers for each. You can also select whether to redirect HTTP requests to HTTPS. Finally, you can specify whether the IP interface is the default interface for the spam quarantine, and whether to use the hostname as the URL or provide a custom URL.

Step 11 Submit and commit your changes.

Accessing the Appliance via FTP

Caution By disabling services via the Management Appliance > Network > IP Interfaces page or the interfaceconfig command, you can disconnect yourself from the GUI or CLI, depending on how you are connected to the appliance. Do not disable services with this command if you are not able to reconnect to the appliance using another protocol, the Serial interface, or the default settings on the Management port.

Step 1 [New Web Interface Only] On the Cloud Email Security Management Console, click on the gear icon to load the legacy web interface.

Step 2 Choose **Management Appliance** > **Network** > **IP Interfaces** page (or the `interfaceconfig` command) to enable FTP access for the interface.

Note Remember to commit your changes before moving on to the next step.

Step 3 Access the interface via FTP. Ensure you are using the correct IP address for the interface.
Example: `ftp 192.168.42.42`

Many browsers also allow you to access interfaces via FTP.

Example: `ftp://192.10.10.10`

Step 4 Browse to the directory for the specific task you are trying to accomplish. After you have accessed an interface via FTP, you can browse the following directories to copy and add ("`GET`" and "`PUT`") files. See the following table.

Table 84: Directories Available for Access

Directory Name	Description

`/avarchive`
`/bounces`
`/cli_logs`
`/delivery`
`/error_logs`
`/ftpd_logs`
`/gui_logs`
`/mail_logs`
`/rptd_logs`
`/sntpd.logs`
`/status`
`/system_logs`
	Created automatically for logging via the Management Appliance > System Administration > Log Subscriptions page or the logconfig and rollovernow commands. See the "Logging" chapter in the user guide or online help for your Email Security appliance for a detailed description of each log.
	See "Log File Type Comparison" in the "Logging" chapter for the differences among each log file type.

`/configuration`	The directory where data from the following pages and commands are exported to and/or imported (saved) from: • Virtual Gateway mappings (`altsrchost`) • Configuration data in XML format (`saveconfig`, `loadconfig`) • Host Access Table (HAT) page (`hostaccess`) • Recipient Access Table (RAT) page (`rcptaccess`) • SMTP Routes page (`smtproutes`) • Alias tables (`aliasconfig`) • Masquerading tables (`masquerade`) • Message filters (`filters`) • Global unsubscribe data (`unsubscribe`) • Test messages for the `trace` command

Directory Name	Description
/MFM	The Mail Flow Monitoring database directory contains data for the Mail Flow Monitor functionality available from the GUI. Each subdirectory contains a README file that documents the record format for each file. You can copy these files to a different machine for record keeping, or load the files into a database and create your own analysis application. The record format is the same for all files in all directories; this format may change in future releases.
/periodic_reports	The directory where all archived reports configured on the system are stored.

Step 5 Use your FTP program to upload and download files to and from the appropriate directory.

Secure Copy (scp) Access

If your client operating system supports a secure copy (scp) command, you can copy files to and from the directories listed in the table *Directories Available for Access*. For example, in the following example, the file /tmp/test.txt is copied from the client machine to the configuration directory of the appliance with the hostname mail3.example.com .

Note The command prompts for the user's password (admin). This example is shown for reference only; your operating system's implementation of secure copy may vary.

```
% scp /tmp/test.txt admin@mail3.example.com:configuration
The authenticity of host 'mail3.example.com (192.168.42.42)' can't be established.
DSA key fingerprint is 69:02:01:1d:9b:eb:eb:80:0c:a1:f5:a6:61:da:c8:db.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'mail3.example.com ' (DSA) to the list of known hosts.
admin@mail3.example.com's password: (type the password)
test.txt             100% |***************************|  1007       00:00
%
```

In this example, the same file is copied from the appliance to the client machine:

```
% scp admin@mail3.example.com:configuration/text.txt .
admin@mail3.example.com's password: (type the password)
test.txt             100% |***************************|  1007       00:00
```

You can use secure copy (scp) as an alternative to FTP to transfer files to and from the content security appliance.

Note Only users in the operators and administrators group can use secure copy (scp) to access the appliance. For more information, see About Reverting to an Earlier Version of AsyncOS, on page 310.

Accessing via a Serial Connection

If you are connecting to the appliance via a serial connection, use the following information for the console port.

Complete information about this port is in the hardware installation guide for your appliance.

Related Topics

- Documentation , on page 405

Pinout Details for the Serial Port in 80- and 90- Series Hardware

Figure 20: Pinout Details for the Serial Port in 80- and 90- Series Hardware

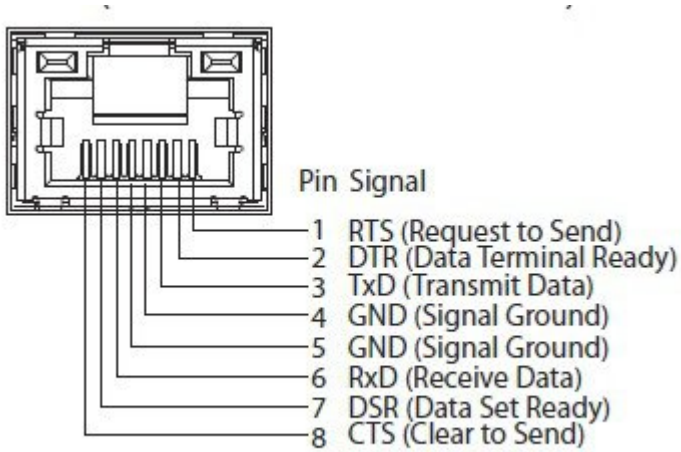

Pinout Details for the Serial Port in 70-Series Hardware

The following figure illustrates the pin numbers for the serial port connector, and the table *Serial Port Pin Assignments* defines the pin assignments and interface signals for the serial port connector.

Figure 21: Pin Numbers for the Serial Port

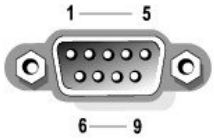

Table 85: Serial Port Pin Assignments

Pin	Signal	I/O	Definition
1	DCD		Data carrier detect
2	SIN		Serial input
3	SOUT		Serial output
4	DTR		Data terminal ready
5	GND	n/a	Signal ground
6	DSR		Data set ready
7	RTS		Request to send
8	CTS		Clear to send
9	RI		Ring indicator
Shell	n/a	n/a	Chassis ground

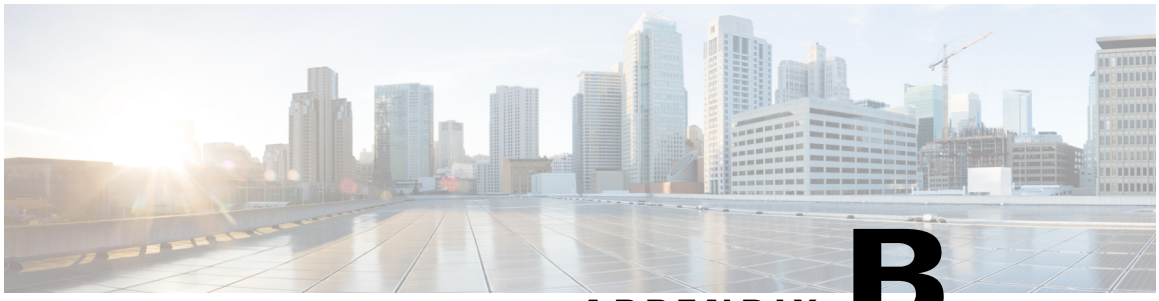

APPENDIX B

Assigning Network and IP Addresses

This appendix contains the following sections:

Ethernet Interfaces

Cisco content security appliances have up to four Ethernet interfaces located on the rear panel of the system, depending on the configuration (whether or not you have the optional optical network interface). They are labeled:

- Management
- Data1
- Data2
- Data3
- Data4

Selecting IP Addresses and Netmasks

When you configure the network, the content security appliance must be able to select a unique interface to send an outgoing packet. This requirement drives some of the decisions regarding IP address and netmask selection for the Ethernet interfaces. The rule is that only one interface can be on a single network (as determined through the applications of netmasks to the IP addresses of the interfaces).

An IP address identifies a physical interface on any given network. A physical Ethernet interface can have more than one IP address for which it accepts packets. An Ethernet interface that has more than one IP address can send packets over that interface with any one of the IP addresses as the source address in the packet. This property is used in implementing Virtual Gateway technology.

The purpose of a netmask is to divide an IP address into a network address and a host address. The network address can be thought of as the network part (the bits matching the netmask) of the IP address. The host address is the remaining bits of the IP address. The number of bits in a four octet address that are significant are sometimes expressed in Classless Inter-Domain Routing (CIDR) style. This is a slash followed by the number of bits (1-32).

A netmask can be expressed in this way by simply counting the ones in binary, so 255.255.255.0 becomes " /24 " and 255.255.240.0 becomes " /20. "

Sample Interface Configurations

This section shows sample interface configurations based on some typical networks. The example uses two interfaces called Int1 and Int2. In the case of the content security appliance, these interface names can represent any two interfaces out of the three interfaces (Management, Data1, Data2).

Network 1:

Separate interfaces must appear to be on separate networks.

Interface	IP Address	Netmask	Net Address
Int1	192.168.1.10	255.255.255.0	192.168.1.0/24
Int2	192.168.0.10	255.255.255.0	192.168.0.0/24

Data addressed to 192.168.1.X (where X is any number from 1 through 255, except for your own address, 10 in this case) go out on Int1. Anything addressed to 192.168.0.X goes out on Int2. Any packet headed for some other address not in these formats, most likely out on a WAN or the Internet, is sent to the default gateway, which must be on one of these networks. The default gateway then forwards the packet on.

Network 2:

The network addresses (network parts of the IP addresses) of two different interfaces cannot be the same.

Ethernet Interface	IP Address	Netmask	Net Address
Int1	192.168.1.10	255.255.0.0	192.168.0.0/16
Int2	192.168.0.10	255.255.0.0	192.168.0.0/16

This situation presents a conflict in that two different Ethernet interfaces have the same network address. If a packet from the content security appliance is sent to 192.168.1.11 , there is no way to decide which Ethernet interface should be used to deliver the packet. If the two Ethernet interfaces are connected to two separate physical networks, the packet may be delivered to the incorrect network and never find its destination. The content security appliance does not allow you to configure your network with conflicts.

You can connect two Ethernet interfaces to the same physical network, but you must construct IP addresses and netmasks to allow the content security appliance to select a unique delivery interface.

IP Addresses, Interfaces, and Routing

When you select an interface on which to perform a command or function in the GUI or CLI that allows you to select an interface (for example, upgrading AsyncOS or configuring DNS), routing (your default gateway) takes precedence over your selection.

For example, suppose that you have a content security appliance with the three network interfaces configured, each on a different network segment (assume all /24):

Ethernet	IP
Management	192.19.0.100
Data1	192.19.1.100
Data2	192.19.2.100

And your default gateway is 192.19.0.1.

Now, if you perform an AsyncOS upgrade (or other command or function that allows you to select an interface) and you select the IP that is on Data1 (192.19.1.100), you would expect all the TCP traffic to occur over the Data1 Ethernet interface. However, instead the traffic goes out of the interface that is set as your default gateway, in this case Management, but is stamped with the source address of the IP on Data1.

Summary

The content security appliance must always be able to identify a unique interface over which a packet can be delivered. To make this decision, the content security appliance uses a combination of the packet's destination IP address, and the network and IP address settings of its Ethernet interfaces. The following table summarizes the preceding examples:

	Same Network	Different Network
Same Physical Interface	Allowed	Allowed
Different Physical Interface	Not allowed	Allowed

Strategies for Connecting Your Content Security Appliance

Keep the following in mind when connecting your appliance:

- Administrative traffic (CLI, web interface, log delivery) is usually little compared to email traffic.
- If two Ethernet interfaces are connected to the same network switch, but end up talking to a single interface on another host downstream, or are connected to a network hub where all data are echoed to all ports, no advantage is gained by using two interfaces.

- SMTP conversations over an interface operating at 1000Base-T are slightly faster than conversations over the same interfaces operating at 100Base-T, but only under ideal conditions.
- There is no point in optimizing connections to your network if there is a bottleneck in some other part of your delivery network. Bottlenecks most often occur in the connection to the Internet and further upstream at your connectivity provider.

The number of interfaces that you choose to connect and how you address them should be dictated by the complexity of your underlying network. It is not necessary to connect multiple interfaces if your network topology or data volumes do not call for it. It is also possible to keep the connection simple at first as you familiarize yourself with the gateway and then increase the connectivity as volume and network topology require it.

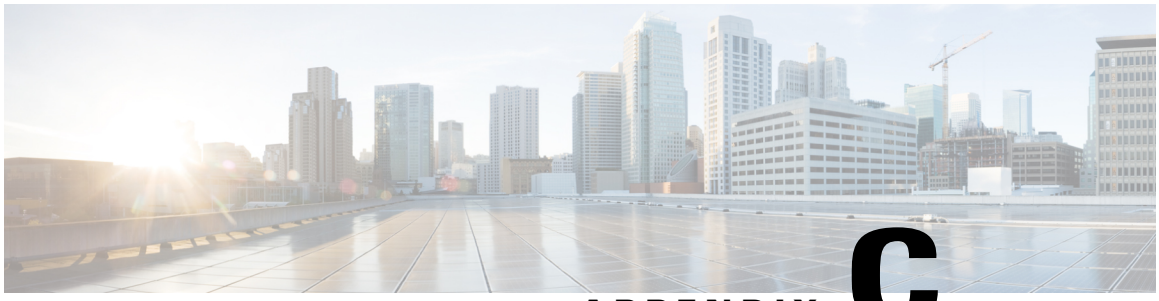

APPENDIX C

Firewall Information

This chapter contains the following sections:

Firewall Information

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

Table 86: Firewall Ports

Default Port	Protocol	In/Out	Hostname	Purpose
20/21	TCP	In or out	AsyncOS IPs, FTP server	FTP for aggregation of log files. Data ports TCP 1024 and higher must also all be open. For more information, search for FTP port information in the Knowledge Base. See Knowledge Base Articles (TechNotes) , on page 407 .
22	SSH	Out	AsyncOS IPs	Centralized configuration manager configuration push. Also used for backups.
22	TCP	In	AsyncOS IPs	SSH access to the CLI, aggregation of log files.
22	TCP	Out	SCP server	SCP push to log server.
23	Telnet	In	AsyncOS IPs	Telnet access to the CLI.

23	Telnet	Out	Telnet Server	Telnet upgrades
25	TCP	Out	Any	SMTP to send email.
25	TCP	In	AsyncOS IPs	SMTP to receive bounced email or if injecting email from outside firewall.
53	UDP/TCP	Out	DNS servers	DNS if configured to use Internet root servers or other DNS servers outside the firewall. Also for SenderBase queries.
80	HTTP	In	AsyncOS IPs	HTTP access to the GUI for system monitoring.
80	HTTP	Out	downloads.ironport.com	Service updates, except for AsyncOS upgrades .
80	HTTP	Out	upgrades.ironport.com	AsyncOS upgrades.
82	HTTP	In	AsyncOS IPs	Used for viewing the spam quarantine.
83	HTTPS	In	AsyncOS IPs	Used for viewing the spam quarantine.
110	TCP	Out	POP server	POP authentication for end users for spam quarantine.
123	UDP	In & Out	NTP server	NTP if time servers are outside firewall.
143	TCP	Out	IMAP server	IMAP authentication for end users for spam quarantine.
161	UDP	In	AsyncOS IPs	SNMP Queries.
162	UDP	Out	Management station	SNMP Traps.
389 or 3268	LDAP	Out	LDAP servers	LDAP if LDAP directory servers are outside firewall. LDAP authentication for Cisco Spam Quarantine.
6363269	LDAPS	Out	LDAPS	LDAPS — ActiveDirectory's global catalog server (uses SSL).
443	TCP	In	AsyncOS IPs	Secure HTTP (https) access to the GUI for system monitoring.

443	TCP	Out	update-static.ironport.com	Verify the latest files for the update server.
443	TCP	Out	update-manifests.ironport.com	Obtain the list of the latest files from the update server (for physical hardware appliances.)
443	TCP	Out	update-manifests.sco.cisco.com	Obtain the list of the latest files from the update server (for virtual appliances.)
443	TCP	Out	phonehome.senderbase.org	Receive/send Outbreak Filters.
443	TCP	Out	File Analysis server URL as configured on your Email Security appliance on the Security Services > File Reputation and Analysis page, in the Advanced Settings for File Analysis section.	Display detailed file analysis results on the File Analysis server.
514	UDP/TCP	Out	Syslog server	Syslog logging.
1024 and higher	—	—	—	See information above for Port 21 (FTP.)
6025	TCP			
7025	TCP	In and out	AsyncOS IPs	Pass policy, virus, and outbreak quarantine data between Email Security appliances and the Security Management appliance when this feature is centralized.
32137	TCP			

Firewall Information

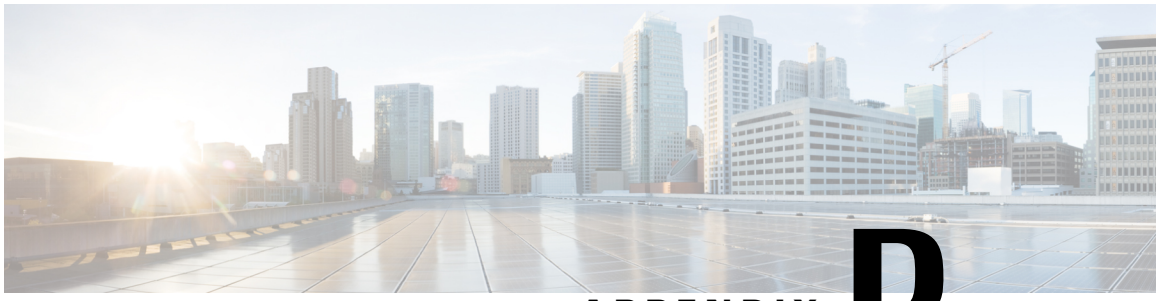

APPENDIX D

Additional Resources

This chapter contains the following sections:

Cisco Notification Service

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

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

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

A Cisco.com account is required. If you do not have one, see Registering for a Cisco Account , on page 407.

Documentation

Documentation for this product and related products is available at the following locations:

Documentation For Cisco Content Security Products:	Is Located At:
Security Management appliances	http://www.cisco.com/c/en/us/support/security/content-security-management-appliance/tsd-products-support-series-home.html Hardware and virtual appliance information: http://www.cisco.com/c/en/us/support/security/content-security-management-appliance/products-installation-guides-list.html MIBs: See Monitoring System Health Using SNMP , on page 288.
Email Security appliances	http://www.cisco.com/c/en/us/support/security/email-security-appliance/tsd-products-support-series-home.html
Command Line Reference guide for content security products	http://www.cisco.com/c/en/us/support/security/email-security-appliance/products-command-reference-list.html
Cisco Email Encryption	http://www.cisco.com/c/en/us/support/security/email-encryption/tsd-products-support-series-home.html

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

Third Party Contributors

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

Information about third-party licenses is available in a Licensing document at: http://www.cisco.com/c/en/us/support/security/content-security-management-appliance/products-release-notes-list.html and at https://support.ironport.com/3rdparty/AsyncOS_User_Guide-1-1.html

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

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

Training

For training options, see:

- http://www.cisco.com/c/en/us/training-events/training-certifications/supplemental-training/email-and-web-security.html
- http://www.cisco.com/c/en/us/training-events/training-certifications/overview.html

Knowledge Base Articles (TechNotes)

DETAILED STEPS

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

Cisco Support Community

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

Access the Cisco Support Community at the following URLs:

- For email security and associated management: https://supportforums.cisco.com/community/5756/email-security
- For web security and associated management: https://supportforums.cisco.com/community/5786/web-security

Customer Support

Use the following methods to obtain support:

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

Support site for legacy IronPort: http://www.cisco.com/web/services/acquisitions/ironport.html

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

See also Opening or Updating a Support Case from the Appliance , on page 383.

For virtual appliances, see the *Cisco Content Security Virtual Appliance Installation Guide.*

Registering for a Cisco Account

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

If you do not have a Cisco.com User ID, you can register for one here: https://tools.cisco.com/RPF/register/register.do

Related Topics

- Cisco Notification Service , on page 405
- Knowledge Base Articles (TechNotes) , on page 407

Cisco Welcomes Your Comments

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

contentsecuritydocs@cisco.com

Please include the title of this book and the publication date from the title page in the subject line of your message.

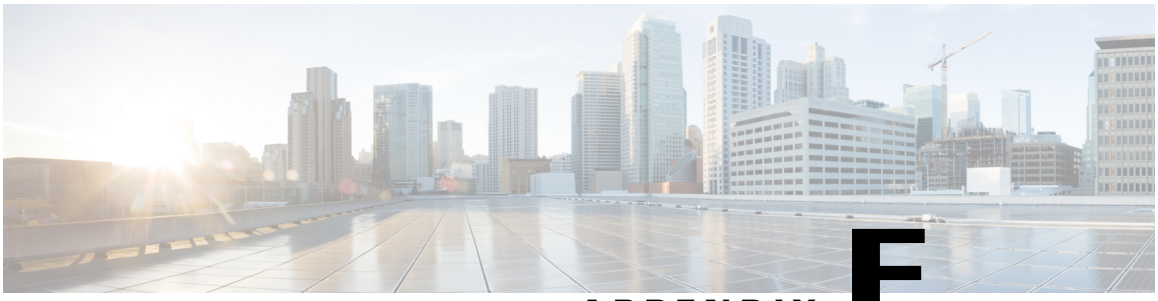

APPENDIX E

End User License Agreement

This chapter contains the following sections:

Cisco Systems End User License Agreement

IMPORTANT: PLEASE READ THIS END USER LICENSE AGREEMENT CAREFULLY. IT IS VERY IMPORTANT THAT YOU CHECK THAT YOU ARE PURCHASING CISCO SOFTWARE OR EQUIPMENT FROM AN APPROVED SOURCE AND THAT YOU, OR THE ENTITY YOU REPRESENT (COLLECTIVELY, THE "CUSTOMER") HAVE BEEN REGISTERED AS THE END USER FOR THE PURPOSES OF THIS CISCO END USER LICENSE AGREEMENT. IF YOU ARE NOT REGISTERED AS THE END USER YOU HAVE NO LICENSE TO USE THE SOFTWARE AND THE LIMITED WARRANTY IN THIS END USER LICENSE AGREEMENT DOES NOT APPLY. ASSUMING YOU HAVE PURCHASED FROM AN APPROVED SOURCE, DOWNLOADING, INSTALLING OR USING CISCO OR CISCO-SUPPLIED SOFTWARE CONSTITUTES ACCEPTANCE OF THIS AGREEMENT.

CISCO SYSTEMS, INC. OR ITS SUBSIDIARY LICENSING THE SOFTWARE INSTEAD OF CISCO SYSTEMS, INC. ("CISCO") IS WILLING TO LICENSE THIS SOFTWARE TO YOU ONLY UPON THE CONDITION THAT YOU PURCHASED THE SOFTWARE FROM AN APPROVED SOURCE AND THAT YOU ACCEPT ALL OF THE TERMS CONTAINED IN THIS END USER LICENSE AGREEMENT PLUS ANY ADDITIONAL LIMITATIONS ON THE LICENSE SET FORTH IN A SUPPLEMENTAL LICENSE AGREEMENT ACCOMPANYING THE PRODUCT OR AVAILABLE AT THE TIME OF YOUR ORDER (COLLECTIVELY THE "AGREEMENT"). TO THE EXTENT OF ANY CONFLICT BETWEEN THE TERMS OF THIS END USER LICENSE AGREEMENT AND ANY SUPPLEMENTAL LICENSE AGREEMENT, THE SUPPLEMENTAL LICENSE AGREEMENT SHALL APPLY. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU ARE REPRESENTING THAT YOU PURCHASED THE SOFTWARE FROM AN APPROVED SOURCE AND BINDING YOURSELF TO THE AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THE AGREEMENT, THEN CISCO IS UNWILLING TO LICENSE THE SOFTWARE TO YOU AND (A) YOU MAY NOT DOWNLOAD, INSTALL OR USE THE SOFTWARE, AND (B) YOU MAY RETURN THE SOFTWARE (INCLUDING ANY UNOPENED CD PACKAGE AND ANY WRITTEN MATERIALS) FOR A FULL REFUND, OR, IF THE SOFTWARE AND WRITTEN MATERIALS ARE SUPPLIED AS PART OF ANOTHER PRODUCT, YOU MAY RETURN THE ENTIRE PRODUCT FOR A FULL REFUND. YOUR

RIGHT TO RETURN AND REFUND EXPIRES 30 DAYS AFTER PURCHASE FROM AN APPROVED SOURCE, AND APPLIES ONLY IF YOU ARE THE ORIGINAL AND REGISTERED END USER PURCHASER. FOR THE PURPOSES OF THIS END USER LICENSE AGREEMENT, AN "APPROVED SOURCE" MEANS (A) CISCO; OR (B) A DISTRIBUTOR OR SYSTEMS INTEGRATOR AUTHORIZED BY CISCO TO DISTRIBUTE / SELL CISCO EQUIPMENT, SOFTWARE AND SERVICES WITHIN YOUR TERRITORY TO END USERS; OR (C) A RESELLER AUTHORIZED BY ANY SUCH DISTRIBUTOR OR SYSTEMS INTEGRATOR IN ACCORDANCE WITH THE TERMS OF THE DISTRIBUTOR'S AGREEMENT WITH CISCO TO DISTRIBUTE / SELL THE CISCO EQUIPMENT, SOFTWARE AND SERVICES WITHIN YOUR TERRITORY TO END USERS.

THE FOLLOWING TERMS OF THE AGREEMENT GOVERN CUSTOMER'S USE OF THE SOFTWARE (DEFINED BELOW), EXCEPT TO THE EXTENT: (A) THERE IS A SEPARATE SIGNED CONTRACT BETWEEN CUSTOMER AND CISCO GOVERNING CUSTOMER'S USE OF THE SOFTWARE, OR (B) THE SOFTWARE INCLUDES A SEPARATE "CLICK-ACCEPT" LICENSE AGREEMENT OR THIRD PARTY LICENSE AGREEMENT AS PART OF THE INSTALLATION OR DOWNLOAD PROCESS GOVERNING CUSTOMER'S USE OF THE SOFTWARE. TO THE EXTENT OF A CONFLICT BETWEEN THE PROVISIONS OF THE FOREGOING DOCUMENTS, THE ORDER OF PRECEDENCE SHALL BE (1)THE SIGNED CONTRACT, (2) THE CLICK-ACCEPT AGREEMENT OR THIRD PARTY LICENSE AGREEMENT, AND (3) THE AGREEMENT. FOR PURPOSES OF THE AGREEMENT, "SOFTWARE" SHALL MEAN COMPUTER PROGRAMS, INCLUDING FIRMWARE AND COMPUTER PROGRAMS EMBEDDED IN CISCO EQUIPMENT, AS PROVIDED TO CUSTOMER BY AN APPROVED SOURCE, AND ANY UPGRADES, UPDATES, BUG FIXES OR MODIFIED VERSIONS THERETO (COLLECTIVELY, "UPGRADES"), ANY OF THE SAME WHICH HAS BEEN RELICENSED UNDER THE CISCO SOFTWARE TRANSFER AND RE-LICENSING POLICY (AS MAY BE AMENDED BY CISCO FROM TIME TO TIME) OR BACKUP COPIES OF ANY OF THE FOREGOING.

License. Conditioned upon compliance with the terms and conditions of the Agreement, Cisco grants to Customer a nonexclusive and nontransferable license to use for Customer's internal business purposes the Software and the Documentation for which Customer has paid the required license fees to an Approved Source. "Documentation" means written information (whether contained in user or technical manuals, training materials, specifications or otherwise) pertaining to the Software and made available by an Approved Source with the Software in any manner (including on CD-Rom, or on-line). In order to use the Software, Customer may be required to input a registration number or product authorization key and register Customer's copy of the Software online at Cisco's website to obtain the necessary license key or license file.

Customer's license to use the Software shall be limited to, and Customer shall not use the Software in excess of, a single hardware chassis or card or such other limitations as are set forth in the applicable Supplemental License Agreement or in the applicable purchase order which has been accepted by an Approved Source and for which Customer has paid to an Approved Source the required license fee (the "Purchase Order").

Unless otherwise expressly provided in the Documentation or any applicable Supplemental License Agreement, Customer shall use the Software solely as embedded in, for execution on, or (where the applicable Documentation permits installation on non-Cisco equipment) for communication with Cisco equipment owned or leased by Customer and used for Customer's internal business purposes. No other licenses are granted by implication, estoppel or otherwise.

For evaluation or beta copies for which Cisco does not charge a license fee, the above requirement to pay license fees does not apply.

General Limitations. This is a license, not a transfer of title, to the Software and Documentation, and Cisco retains ownership of all copies of the Software and Documentation. Customer acknowledges that the Software and Documentation contain trade secrets of Cisco or its suppliers or licensors, including but not limited to the specific internal design and structure of individual programs and associated interface information. Except as otherwise expressly provided under the Agreement, Customer shall only use the Software in connection with

the use of Cisco equipment purchased by the Customer from an Approved Source and Customer shall have no right, and Customer specifically agrees not to:

(i) transfer, assign or sublicense its license rights to any other person or entity (other than in compliance with any Cisco relicensing/transfer policy then in force), or use the Software on Cisco equipment not purchased by the Customer from an Approved Source or on secondhand Cisco equipment, and Customer acknowledges that any attempted transfer, assignment, sublicense or use shall be void;

(ii) make error corrections to or otherwise modify or adapt the Software or create derivative works based upon the Software, or permit third parties to do the same;

(iii) reverse engineer or decompile, decrypt, disassemble or otherwise reduce the Software to human-readable form, except to the extent otherwise expressly permitted under applicable law notwithstanding this restriction or except to the extent that Cisco is legally required to permit such specific activity pursuant to any applicable open source license;

(iv) publish any results of benchmark tests run on the Software;

(v) use or permit the Software to be used to perform services for third parties, whether on a service bureau or time sharing basis or otherwise, without the express written authorization of Cisco; or

(vi) disclose, provide, or otherwise make available trade secrets contained within the Software and Documentation in any form to any third party without the prior written consent of Cisco. Customer shall implement reasonable security measures to protect such trade secrets.

To the extent required by applicable law, and at Customer's written request, Cisco shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of Cisco's applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Cisco makes such information available.

Software, Upgrades and Additional Copies. NOT WITHSTANDING ANY OTHER PROVISION OF THE AGREEMENT: (1) CUSTOMER HAS NO LICENSE OR RIGHT TO MAKE OR USE ANY ADDITIONAL COPIES OR UPGRADES UNLESS CUSTOMER, AT THE TIME OF MAKING OR ACQUIRING SUCH COPY OR UPGRADE, ALREADY HOLDS A VALID LICENSE TO THE ORIGINAL SOFTWARE AND HAS PAID THE APPLICABLE FEE TO AN APPROVED SOURCE FOR THE UPGRADE OR ADDITIONAL COPIES; (2) USE OF UPGRADES IS LIMITED TO CISCO EQUIPMENT SUPPLIED BY AN APPROVED SOURCE FOR WHICH CUSTOMER IS THE ORIGINAL END USER PURCHASER OR LESSEE OR OTHERWISE HOLDS A VALID LICENSE TO USE THE SOFTWARE WHICH IS BEING UPGRADED; AND (3) THE MAKING AND USE OF ADDITIONAL COPIES IS LIMITED TO NECESSARY BACKUP PURPOSES ONLY.

Proprietary Notices. Customer agrees to maintain and reproduce all copyright, proprietary, and other notices on all copies, in any form, of the Software in the same form and manner that such copyright and other proprietary notices are included on the Software. Except as expressly authorized in the Agreement, Customer shall not make any copies or duplicates of any Software without the prior written permission of Cisco.

Term and Termination. The Agreement and the license granted herein shall remain effective until terminated. Customer may terminate the Agreement and the license at any time by destroying all copies of Software and any Documentation. Customer's rights under the Agreement will terminate immediately without notice from Cisco if Customer fails to comply with any provision of the Agreement. Upon termination, Customer shall destroy all copies of Software and Documentation in its possession or control. All confidentiality obligations of Customer, all restrictions and limitations imposed on the Customer under the section titled "General Limitations" and all limitations of liability and disclaimers and restrictions of warranty shall survive termination of this Agreement. In addition, the provisions of the sections titled "U.S. Government End User Purchasers" and "General Terms Applicable to the Limited Warranty Statement and End User License Agreement" shall survive termination of the Agreement.

Customer Records. Customer grants to Cisco and its independent accountants the right to examine Customer's books, records and accounts during Customer's normal business hours to verify compliance with this Agreement. In the event such audit discloses non-compliance with this Agreement, Customer shall promptly pay to Cisco the appropriate license fees, plus the reasonable cost of conducting the audit.

Export, Re-Export, Transfer and Use Controls. The Software, Documentation and technology or direct products thereof (hereafter referred to as Software and Technology), supplied by Cisco under the Agreement are subject to export controls under the laws and regulations of the United States (U.S.) and any other applicable countries' laws and regulations. Customer shall comply with such laws and regulations governing export, re-export, transfer and use of Cisco Software and Technology and will obtain all required U.S. and local authorizations, permits, or licenses. Cisco and Customer each agree to provide the other information, support documents, and assistance as may reasonably be required by the other in connection with securing authorizations or licenses. Information regarding compliance with export, re-export, transfer and use may be located at the following URL:

http://www.cisco.com/c/en/us/about/legal/global-export-trade/general-export/contract-compliance.html.

U.S. Government End User Purchasers. The Software and Documentation qualify as "commercial items," as that term is defined at Federal Acquisition Regulation ("FAR") (48 C.F.R.) 2.101, consisting of "commercial computer software" and "commercial computer software documentation" as such terms are used in FAR 12.212. Consistent with FAR 12.212 and DoD FAR Supp. 227.7202-1 through 227.7202-4, and notwithstanding any other FAR or other contractual clause to the contrary in any agreement into which the Agreement may be incorporated, Customer may provide to Government end user or, if the Agreement is direct, Government end user will acquire, the Software and Documentation with only those rights set forth in the Agreement. Use of either the Software or Documentation or both constitutes agreement by the Government that the Software and Documentation are "commercial computer software" and "commercial computer software documentation," and constitutes acceptance of the rights and restrictions herein.

Identified Components; Additional Terms. The Software may contain or be delivered with one or more components, which may include third-party components, identified by Cisco in the Documentation, readme.txt file, third-party click-accept or elsewhere (e.g. on http://www.cisco.com/) (the "Identified Component(s)") as being subject to different license agreement terms, disclaimers of warranties, limited warranties or other terms and conditions (collectively, "Additional Terms") than those set forth herein. You agree to the applicable Additional Terms for any such Identified Component(s)."

Limited Warranty

Subject to the limitations and conditions set forth herein, Cisco warrants that commencing from the date of shipment to Customer (but in case of resale by an Approved Source other than Cisco, commencing not more than ninety (90) days after original shipment by Cisco), and continuing for a period of the longer of (a) ninety (90) days or (b) the warranty period (if any) expressly set forth as applicable specifically to software in the warranty card accompanying the product of which the Software is a part (the "Product") (if any): (a) the media on which the Software is furnished will be free of defects in materials and workmanship under normal use; and (b) the Software substantially conforms to the Documentation. The date of shipment of a Product by Cisco is set forth on the packaging material in which the Product is shipped. Except for the foregoing, the Software is provided "AS IS". This limited warranty extends only to the Software purchased from an Approved Source by a Customer who is the first registered end user. Customer's sole and exclusive remedy and the entire liability of Cisco and its suppliers under this limited warranty will be (i) replacement of defective media and/or (ii) at Cisco's option, repair, replacement, or refund of the purchase price of the Software, in both cases subject to the condition that any error or defect constituting a breach of this limited warranty is reported to the Approved Source supplying the Software to Customer, within the warranty period. Cisco or the Approved Source supplying the Software to Customer may, at its option, require return of the Software and/or Documentation as a condition to the remedy. In no event does Cisco warrant that the Software is error free or that Customer will be able to operate the Software without problems or interruptions. In addition, due to the continual

development of new techniques for intruding upon and attacking networks, Cisco does not warrant that the Software or any equipment, system or network on which the Software is used will be free of vulnerability to intrusion or attack.

Restrictions. This warranty does not apply if the Software, Product or any other equipment upon which the Software is authorized to be used (a) has been altered, except by Cisco or its authorized representative, (b) has not been installed, operated, repaired, or maintained in accordance with instructions supplied by Cisco, (c) has been subjected to abnormal physical or electrical stress, abnormal environmental conditions, misuse, negligence, or accident; or (d) is licensed for beta, evaluation, testing or demonstration purposes. The Software warranty also does not apply to (e) any temporary Software modules; (f) any Software not posted on Cisco's Software Center; (g) any Software that Cisco expressly provides on an "AS IS" basis on Cisco's Software Center; (h) any Software for which an Approved Source does not receive a license fee; and (i) Software supplied by any third party which is not an Approved Source.

DISCLAIMER OF WARRANTY

EXCEPT AS SPECIFIED IN THIS WARRANTY SECTION, ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS, AND WARRANTIES INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTY OR CONDITION OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, SATISFACTORYQUALITY, NON-INTERFERENCE, ACCURACY OF INFORMATIONAL CONTENT, OR ARISING FROM A COURSE OF DEALING, LAW, USAGE, OR TRADE PRACTICE, ARE HEREBY EXCLUDED TO THE EXTENT ALLOWED BY APPLICABLE LAW AND ARE EXPRESSLY DISCLAIMED BY CISCO, ITS SUPPLIERS AND LICENSORS. TO THE EXTENT THAT ANY OF THE SAME CANNOT BE EXCLUDED, SUCH IMPLIEDCONDITION, REPRESENTATION AND/OR WARRANTY IS LIMITED IN DURATION TO THE EXPRESS WARRANTY PERIOD REFERRED TO IN THE "LIMITED WARRANTY" SECTION ABOVE. BECAUSE SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATIONS ON HOW LONG AN IMPLIED WARRANTY LASTS, THE ABOVE LIMITATION MAY NOT APPLY IN SUCH STATES. THIS WARRANTY GIVES CUSTOMER SPECIFIC LEGAL RIGHTS, AND CUSTOMER MAY ALSO HAVE OTHER RIGHTS WHICH VARY FROM JURISDICTION TO JURISDICTION. This disclaimer and exclusion shall apply even if the express warranty set forth above fails of its essential purpose.

Disclaimer of Liabilities - Limitation of Liability. IF YOU ACQUIRED THE SOFTWARE IN THE UNITED STATES, LATIN AMERICA, CANADA, JAPAN OR THE CARIBBEAN, NOTWITHSTANDING ANYTHING ELSE IN THE AGREEMENT TO THE CONTRARY, ALL LIABILITY OF CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS COLLECTIVELY, TO CUSTOMER, WHETHER IN CONTRACT, TORT (INCLUDING NEGLIGENCE), BREACH OF WARRANTY OR OTHERWISE, SHALL NOT EXCEED THE PRICE PAID BY CUSTOMER TO ANY APPROVED SOURCE FOR THE SOFTWARE THAT GAVE RISE TO THE CLAIM OR IF THE SOFTWARE IS PART OF ANOTHER PRODUCT, THE PRICE PAID FOR SUCH OTHER PRODUCT. THIS LIMITATION OF LIABILITY FOR SOFTWARE IS CUMULATIVE AND NOT PER INCIDENT (I.E. THE EXISTENCE OF TWO OR MORE CLAIMS WILL NOT ENLARGE THIS LIMIT).

IF YOU ACQUIRED THE SOFTWARE IN EUROPE, THE MIDDLE EAST, AFRICA, ASIA OR OCEANIA, NOTWITHSTANDING ANYTHING ELSE IN THE AGREEMENT TO THE CONTRARY, ALL LIABILITY OF CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS COLLECTIVELY, TO CUSTOMER, WHETHER IN CONTRACT, TORT (INCLUDING NEGLIGENCE), BREACH OF WARRANTY OR OTHERWISE, SHALL NOT EXCEED THE PRICE PAID BY CUSTOMER TO CISCO FOR THE SOFTWARE THAT GAVE RISE TO THE CLAIM OR IF THE SOFTWARE IS PART OF ANOTHER PRODUCT, THE PRICE PAID FOR SUCH OTHER PRODUCT. THIS LIMITATION OF LIABILITY FOR SOFTWARE IS CUMULATIVE AND NOT PER INCIDENT (I.E. THE EXISTENCE OF TWO OR MORE CLAIMS WILL NOT ENLARGE THIS LIMIT). NOTHING

IN THE AGREEMENT SHALL LIMIT (I) THE LIABILITY OF CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS TO CUSTOMER FOR PERSONAL INJURY OR DEATH CAUSED BY THEIR NEGLIGENCE, (II) CISCO'S LIABILITY FOR FRAUDULENT MISREPRESENTATION, OR (III) ANY LIABILITY OF CISCO WHICH CANNOT BE EXCLUDED UNDER APPLICABLE LAW.

Disclaimer of Liabilities - Waiver of Consequential Damages and Other Losses. IF YOU ACQUIRED THE SOFTWARE IN THE UNITED STATES, LATIN AMERICA, THE CARIBBEAN OR CANADA, REGARDLESS OF WHETHER ANY REMEDY SET FORTH HEREIN FAILS OF ITS ESSENTIAL PURPOSE OR OTHERWISE, IN NO EVENT WILL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY LOST REVENUE, PROFIT, OR LOST OR DAMAGED DATA, BUSINESS INTERRUPTION, LOSS OF CAPITAL, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY OR WHETHER ARISING OUT OF THE USE OF OR INABILITY TO USE SOFTWARE OR OTHERWISE AND EVEN IF CISCO OR ITS SUPPLIERS OR LICENSORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATION OR EXCLUSION OF CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATION MAY NOT APPLY TO YOU.

IF YOU ACQUIRED THE SOFTWARE IN JAPAN, EXCEPT FOR LIABILITY ARISING OUT OF OR IN CONNECTION WITH DEATH OR PERSONAL INJURY, FRAUDULENT MISREPRESENTATION, AND REGARDLESS OF WHETHER ANY REMEDY SET FORTH HEREIN FAILS OF ITS ESSENTIAL PURPOSE OR OTHERWISE, IN NO EVENT WILL CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS BE LIABLE FOR ANY LOST REVENUE, PROFIT, OR LOST OR DAMAGED DATA, BUSINESS INTERRUPTION, LOSS OF CAPITAL, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY OR WHETHER ARISING OUT OF THE USE OF OR INABILITY TO USE SOFTWARE OR OTHERWISE AND EVEN IF CISCO OR ANY APPROVED SOURCE OR THEIR SUPPLIERS OR LICENSORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

IF YOU ACQUIRED THE SOFTWARE IN EUROPE, THE MIDDLE EAST, AFRICA, ASIA OR OCEANIA, IN NO EVENT WILL CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS, BE LIABLE FOR ANY LOST REVENUE, LOST PROFIT, OR LOST OR DAMAGED DATA, BUSINESS INTERRUPTION, LOSS OF CAPITAL, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES, HOWSOEVER ARISING, INCLUDING, WITHOUT LIMITATION, IN CONTRACT, TORT (INCLUDING NEGLIGENCE) OR WHETHER ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF, IN EACH CASE, CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS, HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATION OR EXCLUSION OF CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATION MAY NOT FULLY APPLY TO YOU. THE FOREGOING EXCLUSION SHALL NOT APPLY TO ANY LIABILITY ARISING OUT OF OR IN CONNECTION WITH: (I) DEATH OR PERSONAL INJURY, (II) FRAUDULENT MISREPRESENTATION, OR (III) CISCO'S LIABILITY IN CONNECTION WITH ANY TERMS THAT CANNOT BE EXCLUDED UNDER APPLICABLE LAW.

Customer acknowledges and agrees that Cisco has set its prices and entered into the Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the parties.

Controlling Law, Jurisdiction. If you acquired, by reference to the address on the purchase order accepted by the Approved Source, the Software in the United States, Latin America, or the Caribbean, the Agreement and

warranties ("Warranties") are controlled by and construed under the laws of the State of California, United States of America, notwithstanding any conflicts of law provisions; and the state and federal courts of California shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in Canada, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of the Province of Ontario, Canada, notwithstanding any conflicts of law provisions; and the courts of the Province of Ontario shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in Europe, the Middle East, Africa, Asia or Oceania (excluding Australia), unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of England, notwithstanding any conflicts of law provisions; and the English courts shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. In addition, if the Agreement is controlled by the laws of England, no person who is not a party to the Agreement shall be entitled to enforce or take the benefit of any of its terms under the Contracts (Rights of Third Parties) Act 1999. If you acquired the Software in Japan, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of Japan, notwithstanding any conflicts of law provisions; and the Tokyo District Court of Japan shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in Australia, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of the State of New South Wales, Australia, notwithstanding any conflicts of law provisions; and the State and federal courts of New South Wales shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in any other country, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of the State of California, United States of America, notwithstanding any conflicts of law provisions; and the state and federal courts of California shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties.

For all countries referred to above, the parties specifically disclaim the application of the UN Convention on Contracts for the International Sale of Goods. Notwithstanding the foregoing, either party may seek interim injunctive relief in any court of appropriate jurisdiction with respect to any alleged breach of such party's intellectual property or proprietary rights. If any portion hereof is found to be void or unenforceable, the remaining provisions of the Agreement and Warranties shall remain in full force and effect. Except as expressly provided herein, the Agreement constitutes the entire agreement between the parties with respect to the license of the Software and Documentation and supersedes any conflicting or additional terms contained in any Purchase Order or elsewhere, all of which terms are excluded. The Agreement has been written in the English language, and the parties agree that the English version will govern.

Product warranty terms and other information applicable to Cisco products are available at the following URL:

http://www.cisco.com/c/en/us/products/warranty-listing.html

Supplemental End User License Agreement for Cisco Systems Content Security Software

IMPORTANT: READ CAREFULLY

This Supplemental End User License Agreement ("SEULA") contains additional terms and conditions for the Software product licensed under the End User License Agreement ("EULA") between You ("You" as used herein means You and the business entity you represent or "Company") and Cisco (collectively, the "Agreement"). Capitalized terms used in this SEULA but not defined will have the meanings assigned to them in the EULA. To the extent that there is a conflict between the terms and conditions of the EULA and this SEULA, the terms and conditions of this SEULA will take precedence.

In addition to the limitations set forth in the EULA on your access and use of the Software, you agree to comply at all times with the terms and conditions provided in this SEULA.

DOWNLOADING, INSTALLING, OR USING THE SOFTWARE CONSTITUTES ACCEPTANCE OF THE AGREEMENT, AND YOU ARE BINDING YOURSELF AND THE BUSINESS ENTITY THAT YOU REPRESENT TO THE AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THE AGREEMENT, THEN CISCO IS UNWILLING TO LICENSE THE SOFTWARE TO YOU AND (A) YOU MAY NOT DOWNLOAD, INSTALL OR USE THE SOFTWARE, AND (B) YOU MAY RETURN THE SOFTWARE (INCLUDING ANY UNOPENED CD PACKAGE AND ANY WRITTEN MATERIALS) FOR A FULL REFUND, OR, IF THE SOFTWARE AND WRITTEN MATERIALS ARE SUPPLIED AS PART OF ANOTHER PRODUCT, YOU MAY RETURN THE ENTIRE PRODUCT FOR A FULL REFUND. YOUR RIGHT TO RETURN AND REFUND EXPIRES 30 DAYS AFTER PURCHASE FROM CISCO OR AN AUTHORIZED CISCO RESELLER, AND APPLIES ONLY IF YOU ARE THE ORIGINAL END USER PURCHASER.

For purposes of this SEULA, the Product name and the Product description You have ordered is any of the following Cisco Systems Email Security Appliance ("ESA"), Cisco Systems Web Security Appliance ("WSA") and Cisco Systems Security Management Application ("SMA") (collectively, "Content Security") and their Virtual Appliance equivalent ("Software"):

Cisco AsyncOS for Email

Cisco AsyncOS for Web

Cisco AsyncOS for Management

Cisco Email Anti-Spam, Sophos Anti-Virus

Cisco Email Outbreak Filters

Cloudmark Anti-Spam

Cisco Image Analyzer

McAfee Anti-Virus

Cisco Intelligent Multi-Scan

Cisco RSA Data Loss Prevention

Cisco Email Encryption

Cisco Email Delivery Mode

Cisco Web Usage Controls

Cisco Web Reputation

Sophos Anti-Malware

Webroot Anti-Malware

McAfee Anti-Malware

Cisco Email Reporting

Cisco Email Message Tracking

Cisco Email Centralized Quarantine

Cisco Web Reporting

Cisco Web Policy and Configuration Management

Cisco Advanced Web Security Management with Splunk

Email Encryption for Encryption Appliances

Email Encryption for System Generated Bulk Email

Email Encryption and Public Key Encryption for Encryption Appliances

Large Attachment Handling for Encryption Appliances

Secure Mailbox License for Encryption Appliances

Definitions

For purposes of this SEULA, the following definitions apply:

"Company Service" means the Company's email, Internet, security management services provided to End Users for the purposes of conducting Company's internal business.

"End User" means: (1) for the WSA and SMA, the employee, contractor or other agent authorized by Company to access the Internet and the SMA via the Company Service; and (2) for the ESA, the email boxes of the employees, contractors, or other agent authorized by Company to access or use the email services via the Company Service.

"Ordering Document" means the purchase agreement, evaluation agreement, beta, pre-release agreement or similar agreement between the Company and Cisco or the Company and a Cisco reseller, or the valid terms of any purchase order accepted by Cisco in connection therewith, containing the purchase terms for the Software license granted by this Agreement.

"Personally Identifiable Information" means any information that can be used to identify an individual, including, but not limited to, an individual's name, user name, email address and any other personally identifiable information.

"Server" means a single physical computer or devices on a network that manages or provides network resources for multiple users.

"Services" means Cisco Software Subscription Services.

"Service Description" means the description of the Software Subscription Support Services at http://www.cisco.com/c/en/us/about/legal/service-descriptions.html.

"Telemetry Data" means samples of Company's email and web traffic, including data on email message and web request attributes and information on how different types of email messages and web requests were handled by Company's Cisco hardware products. Email message metadata and web requests included in Telemetry Data are anonymized and obfuscated to remove any Personally Identifiable Information.

"Term" means the length of the Software subscription You purchased, as indicated in your Ordering Document.

"Virtual Appliance" means the virtual version of Cisco's email security appliances, web security appliances, and security management appliances.

"Virtual Machine" means a software container that can run its own operating system and execute applications like a Server.

Additional License Terms and Conditions

LICENSE GRANTS AND CONSENT TO TERMS OF DATA COLLECTION

License of Software.

By using the Software and the Documentation, Company agrees to be bound by the terms of this Agreement, and so long as Company is in compliance with this Agreement, Cisco hereby grants to Company a nonexclusive, non-sublicensable, non-transferable, worldwide license during the Term to use the Software only on Cisco's hardware products, or in the case of the Virtual Appliances, on a Virtual Machine, solely in connection with

the provision of the Company Service to End Users. The number of End Users licensed for the use of the Software is limited to the number of End Users specified in the Ordering Documents. In the event that the number of End Users in connection with the provision of the Company Service exceeds the number of End Users specified in the Ordering Documents, Company shall contact an Approved Source to purchase additional licenses for the Software. The duration and scope of this license(s) is further defined in the Ordering Document. The Ordering Document supersedes the EULA with respect to the term of the Software license. Except for the license rights granted herein, no right, title or interest in any Software is granted to the Company by Cisco, Cisco's resellers or their respective licensors. Your entitlement to Upgrades to the Software is subject to the Service Description. This Agreement and the Services are co-terminus.

Consent and License to Use Data.

Subject to the Cisco Privacy Statement at http://www.cisco.com/web/siteassets/legal/privacy.html, Company hereby consents and grants to Cisco a license to collect and use Telemetry Data from the Company. Cisco does not collect or use Personally Identifiable Information in the Telemetry Data. Cisco may share aggregated and anonymous Telemetry Data with third parties to assist us in improving your user experience and the Software and other Cisco security products and services. Company may terminate Cisco's right to collect Telemetry Data at any time by disabling SenderBase Network Participation in the Software. Instructions to enable or disable SenderBase Network Participation are available in the Software configuration guide.

Description of Other Rights and Obligations

Please refer to the Cisco Systems, Inc. End User License Agreement, Privacy Statement and Service Description of Software Subscription Support Services.

INDEX

Q

R

S

T

U

V

W