Cisco IronPort AsyncOS 8.0 for Security Management User Guide

August 9, 2013

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

Text Part Number: OL-29365-02

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

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

Cisco IronPort AsyncOS 8.0 for Security Management User Guide

© 2008-2013 Cisco Systems, Inc. All rights reserved.

CONTENTS

CHAPTER 1 **Getting Started** 1-1

What's New in This Release 1-1

Where to Find More Information 1-2

Documentation 1-2

Third Party Contributors 1-3

Training 1-3

Knowledge Base 1-3

Cisco Support Community 1-4

Customer Support 1-4

Cisco Welcomes Your Comments 1-4

Security Management Appliance Overview 1-5

CHAPTER 2 **Setup, Installation, and Basic Configuration** 2-1

Solution Deployment Overview 2-1

SMA Compatibility Matrix 2-2

Installation Planning 2-2

Network Planning 2-2

About Integrating a Security Management Appliance with Email Security Appliances 2-3

Deployments with Centralized Management of Email Security Appliances 2-3

Preparing for Setup 2-4

Physically Setting Up and Connecting the Appliance 2-4

Determining Network and IP Address Assignments 2-4

Gathering the Setup Information 2-5

Accessing the Security Management Appliance 2-6

Browser Requirements 2-6

Accessing the Web Interface 2-6

About Accessing the Web Interfaces 2-6

Accessing the Security Management Appliance Command Line Interface 2-7

Supported Languages 2-7

Running the System Setup Wizard 2-8

Before You Begin 2-8

Overview of the System Setup Wizard 2-8

Launch the System Setup Wizard 2-9

Review the End User License Agreement 2-9

Configure the System Settings 2-9
Configure the Network Settings 2-10
Review Your Configuration 2-11
Proceeding to the Next Steps 2-12
About Adding Managed Appliances 2-12
Editing Managed Appliance Configurations 2-13
Removing an Appliance from the List of Managed Appliances 2-14
Configuring Services on the Security Management Appliance 2-14
Committing and Abandoning Configuration Changes 2-14

CHAPTER 3 **Working With Reports** 3-1
Ways to View Reporting Data 3-1
How the Security Appliance Gathers Data for Reports 3-2
How Reporting Data is Stored 3-2
About Reporting and Upgrades 3-3
Customizing Your View of Report Data 3-3
Viewing Reporting Data for an Appliance or Reporting Group 3-4
Choosing a Time Range for Reports 3-4
(Web Reports Only) Choosing Which Data to Chart 3-5
Customizing Tables on Report Pages 3-6
Viewing Details of Transactions Included in Reports 3-6
Improving Performance of Email Reports 3-7
Printing and Exporting Reporting and Tracking Data 3-8
Exporting Report Data as a Comma-Separated Values (CSV) File 3-9
Subdomains vs. Second-Level Domains in Reporting and Tracking 3-10
Email and Web Reports 3-10

CHAPTER 4 **Using Centralized Email Security Reporting** 4-1
Centralized Email Reporting Overview 4-1
Setting Up Centralized Email Reporting 4-2
Enabling Centralized Email Reporting on the Security Management Appliance 4-2
Adding the Centralized Email Reporting Service to Each Managed Email Security Appliance 4-3
Creating Email Reporting Groups 4-4
Enabling Centralized Email Reporting on Email Security Appliances 4-4
Working with Email Report Data 4-4
Searching and the Interactive Email Report Pages 4-5
Understanding the Email Reporting Pages 4-6
Table Column Descriptions for Email Reporting Pages 4-8

Email Reporting Overview Page 4-10
How Incoming Mail Messages are Counted 4-12
How Email Messages Are Categorized by the Appliances 4-12
Categorizing Email Messages on Overview Page 4-13
Incoming Mail Page 4-14
Views Within the Incoming Mail Page 4-14
Categorizing Email Messages on Incoming Mail Page 4-15
Incoming Mail Details Table 4-19
Sender Profile Pages 4-19
Sender Groups Report Page 4-23
Outgoing Destinations Page 4-24
Outgoing Senders Page 4-26
Internal Users Page 4-27
Internal User Details Page 4-29
Searching for a Specific Internal User 4-29
DLP Incident Summary Page 4-30
DLP Incidents Details Table 4-32
DLP Policy Detail Page 4-32
Content Filters Page 4-32
Content Filter Details Page 4-33
Virus Types Page 4-34
TLS Connections Page 4-36
Inbound SMTP Authentication Page 4-39
Rate Limits Page 4-40
Outbreak Filters Page 4-41
System Capacity Page 4-43
How to Interpret the Data You See on System Capacity Page 4-44
System Capacity – Workqueue 4-44
System Capacity – Incoming Mail 4-45
System Capacity – Outgoing Mail 4-46
System Capacity – System Load 4-47
Note About Memory Page Swapping 4-48
System Capacity – All 4-49
Reporting Data Availability Page 4-49
About Scheduled and On-Demand Email Reports 4-50
Additional Report Types 4-52
Domain-Based Executive Summary Report 4-52
Executive Summary Reports 4-55
Scheduling Email Reports 4-55
Adding Scheduled Reports 4-55

Editing Scheduled Reports 4-56
Discontinuing Scheduled Reports 4-57
Generating Email Reports On Demand 4-57
Viewing and Managing Archived Email Reports 4-58
Accessing Archived Reports 4-58
Deleting Archived Reports 4-59

CHAPTER 5 **Using Centralized Web Reporting and Tracking** 5-1
Centralized Web Reporting Overview 5-1
Setting Up Centralized Web Reporting 5-2
Enabling Centralized Web Reporting on the Security Management Appliance 5-3
Enabling Centralized Web Reporting on Web Security Appliances 5-3
Adding the Centralized Web Reporting Service to Each Managed Web Security Appliance 5-3
Anonymizing User Names in Web Reports 5-4
Working with Interactive Web Reporting Pages 5-7
Understanding the Web Reporting Pages 5-7
Table Column Descriptions for Web Reports 5-10
Web Reporting Overview 5-12
Users Report (Web) 5-16
User Details (Web Reporting) 5-19
Web Sites Report 5-22
URL Categories Report 5-24
URL Category Set Updates and Reports 5-26
Using The URL Categories Page in Conjunction with Other Reporting Pages 5-27
Reporting Misclassified and Uncategorized URLs 5-27
Application Visibility Report 5-27
Understanding the Difference between Application versus Application Types 5-28
Anti-Malware Report 5-31
Malware Category Report 5-33
Malware Threat Report 5-34
Malware Category Descriptions 5-35
Client Malware Risk Report 5-37
Web Reputation Filters Report 5-39
What are Web Reputation Filters? 5-39
Adjusting Web Reputation Settings 5-42
L4 Traffic Monitor Report 5-42
SOCKS Proxy Report 5-46
Reports by User Location 5-48
Web Tracking 5-50

Searching for Transactions Processed by Web Proxy Services 5-51
Understanding Web Tracking Search Results 5-53
Searching for Transactions Processed by the L4 Traffic Monitor 5-54
Searching for Transactions Processed by the SOCKS Proxy 5-55
About Web Tracking and Upgrades 5-55
System Capacity Page 5-55
How to Interpret the Data You See on the System Capacity Page 5-56
System Capacity - System Load 5-57
System Capacity - Network Load 5-59
Note About Proxy Buffer Memory Swapping 5-59
Data Availability Page 5-60
About Scheduled and On-Demand Web Reports 5-61
Scheduling Web Reports 5-61
Adding Scheduled Reports 5-62
Editing Scheduled Reports 5-63
Deleting Scheduled Reports 5-63
Additional Extended Reports 5-63
Top URL Categories—Extended 5-63
Top Application Types—Extended 5-64
Generating Web Reports on Demand 5-65
Viewing and Managing Archived Web Reports 5-66

CHAPTER 6 **Tracking Email Messages** 6-1
Tracking Service Overview 6-1
Setting Up Centralized Message Tracking 6-2
Enabling Centralized Email Tracking on a Security Management Appliance 6-2
Configuring Centralized Message Tracking on Email Security Appliances 6-2
Adding the Centralized Message Tracking Service to Each Managed Email Security Appliance 6-3
Managing Access to Sensitive Information 6-4
Searching for Email Messages 6-4
Narrowing the Result Set 6-6
Understanding Tracking Query Results 6-7
Message Details 6-7
Envelope and Header Summary 6-8
Sending Host Summary 6-8
Processing Details 6-8
DLP Matched Content Tab 6-8

CHAPTER 7 **Managing the Cisco IronPort Spam Quarantine** 7-1

Understanding the Cisco IronPort Spam Quarantine 7-1

Setting Up the Centralized Spam Quarantine 7-2

Identifying Required IP Addresses 7-2

Configuring the Cisco IronPort Spam Quarantine Service on the Security Management Appliance 7-2

Configuring Interfaces on the Security Management Appliance 7-4

Configuring an Outbound IP Interface on the Security Management Appliance 7-4

Configuring the IP Interface for Spam Quarantine Access 7-4

Configuring Email Security Appliances for Centralized Spam Quarantine 7-5

Configuring the Email Security Appliance for Centralized Spam Quarantine 7-5

Adding the Centralized Spam Quarantine Service to Each Managed Email Security Appliance 7-6

Configuring Administrative User Access to the Cisco IronPort Spam Quarantine 7-7

Configuring Spam Management Features for End Users 7-8

Configuring End User Quarantine Access 7-8

Configuring Spam Notifications for End Users 7-9

Configuring and Managing the End User Safelist/Blocklist Feature 7-11

Enabling and Configuring Safelist/Blocklists on the Security Management Appliance 7-11

Configuring Safelist/Blocklist Settings on the Email Security Appliance 7-12

Synchronizing Safelist and Blocklist Settings and Databases 7-12

Message Delivery for Safelists and Blocklists 7-13

Backing Up and Restoring the Safelist/Blocklist Database 7-13

Troubleshooting Safelists and Blocklists 7-14

Using End User Safelists and Blocklists 7-14

Accessing Safelists and Blocklists 7-14

Adding Entries to Safelists and Blocklists 7-15

Working with Safelists 7-15

Working with Blocklists 7-16

Managing Messages in the Cisco IronPort Spam Quarantine 7-16

Searching for Messages in the Cisco IronPort Spam Quarantine 7-17

Searching Large Message Collections 7-17

Viewing Messages in the Cisco IronPort Spam Quarantine 7-18

Viewing HTML Messages 7-18

Viewing Encoded Messages 7-18

Delivering Messages in the Cisco IronPort Spam Quarantine 7-18

Deleting Messages from the Cisco IronPort Spam Quarantine 7-18

CHAPTER 8 **Managing Web Security Appliances** 8-1

About Centralized Configuration Management 8-1

Determining the Correct Configuration Publishing Method 8-1
Setting Up Configuration Masters 8-2
Overview of Setting Up Configuration Masters 8-2
Important Notes About Using Configuration Masters 8-3
Determine the Configuration Master Version(s) to Use 8-3
Enabling Centralized Configuration Management on the Security Management Appliance 8-4
Initializing Configuration Masters 8-4
About Associating Web Security Appliances to Configuration Masters 8-5
Adding Web Security Appliances and Associating Them with Configuration Master Versions 8-5
Associating Configuration Master Versions to Web Security Appliances 8-6
Configuring Settings to Publish 8-6
Importing from an Existing Configuration Master 8-7
Importing Settings from a Web Security Appliance 8-7
Configuring Web Security Features Directly in Configuration Masters 8-8
Ensuring that Features are Enabled Consistently 8-10
Comparing Enabled Features 8-10
Enabling Features to Publish 8-11
Disabling Unused Configuration Masters 8-12
Setting Up to Use Advanced File Publishing 8-12
Publishing Configurations to Web Security Appliances 8-13
Publishing a Configuration Master 8-13
Before You Publish a Configuration Master 8-13
Publishing a Configuration Master Now 8-14
Publishing a Configuration Master Later 8-15
Publishing a Configuration Master Using the Command Line Interface 8-16
Publishing Configurations Using Advanced File Publishing 8-16
Advanced File Publish: Publish Configuration Now 8-16
Advanced File Publish: Publish Later 8-17
Viewing Status and History of Publishing Jobs 8-18
Viewing Scheduled Publishing Jobs 8-18
Viewing Status of the Current Publishing Job 8-18
Viewing Publish History 8-18
Viewing Web Security Appliance Status 8-19
Web Appliance Status Page 8-19
Appliance Status Page 8-19
URL Category Set Updates and Centralized Configuration Management 8-22
Understand the Impacts of URL Category Set Updates 8-22
Ensure that You Will Receive Alerts about URL Category Set Updates 8-23

Be Aware: Before You Set Up Configuration Master 7.5 or 7.7 8-23
Specify Default Settings for New and Changed Categories 8-23
When the URL Category Set is Updated, Check Your Policy and Identity Settings 8-23

CHAPTER 9 **Monitoring System Status** 9-1

Monitoring Security Management Appliance Status 9-1
Centralized Services 9-2
Security Appliance Data Transfer Status 9-4
System Information 9-5
Viewing the Status of Your Managed Appliances 9-6
Additional Status Information for Web Security Appliances 9-7
Monitoring Capacity of Managed Appliances 9-7
Monitoring Reporting Data Availability Status 9-7
Monitoring Data Availability on your Email Security Appliances 9-7
Monitoring Data Availability on your Web Security Appliances 9-8
Monitoring Tracking Data Status 9-8
Monitoring Email Tracking Data Status 9-9
Monitoring Web Tracking Data Status 9-9
Identifying Active TCP/IP Services 9-10

CHAPTER 10 **Integrating with LDAP** 10-1

Overview 10-1
Configuring LDAP to Work with the Cisco IronPort Spam Quarantine 10-1
Creating the LDAP Server Profile 10-2
Testing LDAP Servers 10-4
Configuring LDAP Queries 10-4
LDAP Query Syntax 10-4
Tokens 10-5
Spam Quarantine End-User Authentication Queries 10-5
Sample Active Directory End-User Authentication Settings 10-6
Sample OpenLDAP End-User Authentication Settings 10-6
Spam Quarantine Alias Consolidation Queries 10-6
Sample Active Directory Alias Consolidation Settings 10-7
Sample OpenLDAP Alias Consolidation Settings 10-7
Testing LDAP Queries 10-8
Domain-Based Queries 10-8
Creating a Domain-Based Query 10-9
Chain Queries 10-10

Creating a Chain Query 10-10
Configuring AsyncOS to Work With Multiple LDAP Servers 10-11
Testing Servers and Queries 10-12
Failover 10-12
Configuring the Cisco IronPort Appliance for LDAP Failover 10-12
Load Balancing 10-13
Configuring the Cisco IronPort Appliance for Load Balancing 10-13
Configuring External Authentication of Administrative Users Using LDAP 10-14
User Accounts Query for Authenticating Administrative Users 10-15
Group Membership Queries for Authenticating Administrative Users 10-15
Enabling External Authentication of Administrative Users 10-17

CHAPTER 11 **Configuring SMTP Routing** 11-1
Routing Email for Local Domains 11-1
SMTP Routes Overview 11-1
Default SMTP Route 11-2
Defining an SMTP Route 11-2
SMTP Routes Limits 11-3
SMTP Routes and DNS 11-3
SMTP Routes, Mail Delivery, and Message Splintering 11-3
SMTP Routes and Outbound SMTP Authentication 11-3
Managing SMTP Routes on the Security Management Appliance 11-3
Adding SMTP Routes 11-4
Editing SMTP Routes 11-4
Deleting SMTP Routes 11-4
Exporting SMTP Routes 11-4
Importing SMTP Routes 11-5

CHAPTER 12 **Distributing Administrative Tasks** 12-1
About Distributing Administrative Tasks 12-1
Assigning User Roles 12-1
Predefined User Roles 12-2
Custom User Roles 12-4
About Custom Email User Roles 12-4
Creating Custom Email User Roles 12-6
Using Custom Email User Roles 12-7
About Custom Web User Roles 12-7
Creating Custom Web User Roles 12-8
Editing Custom Web User Roles 12-9

Managing Authentication of Administrative Users 12-9
Managing Locally-Defined Administrative Users 12-9
Adding Locally-Defined Users 12-10
Editing Locally-Defined Users 12-10
Deleting Locally-Defined Users 12-11
Viewing the List of Locally-Defined Users 12-11
Setting and Changing Passwords 12-12
Configuring Restrictive User Account and Password Settings 12-12
Requiring Users to Change Passwords at Next Login 12-14
Locking and Unlocking Local User Accounts 12-15
External User Authentication 12-16
Configuring LDAP Authentication 12-16
Enabling RADIUS Authentication 12-16
Additional Controls on Access to the Security Management Appliance 12-18
Configuring IP-Based Network Access 12-18
Direct Connections 12-18
Connecting Through a Proxy 12-18
Creating the Access List 12-19
Configuring the Web UI Session Timeout 12-20
Controlling Access to Sensitive DLP Information in Message Tracking 12-21
Viewing Administrative User Activity 12-21
Viewing Active Sessions Using the Web 12-21
Viewing Administrative User Activity via the Command Line Interface 12-22

CHAPTER 13 **Common Administrative Tasks** 13-1
Working with Feature Keys 13-2
Feature Keys Page 13-2
Feature Key Settings Page 13-3
Expired Feature Keys 13-3
Performing Maintenance Tasks Using CLI Commands 13-3
Shutting Down the Security Management Appliance 13-3
Rebooting the Security Management Appliance 13-3
Placing the Security Management Appliance into a Maintenance State 13-4
The suspend and offline Commands 13-4
Resuming from an Offline State 13-4
The resume Command 13-5
Resetting to Factory Defaults 13-5
The resetconfig Command 13-5
Displaying the Version Information for AsyncOS 13-6

Backing Up Your Security Management Appliance 13-6
What Data Is Backed Up 13-6
Restrictions and Requirements for Backups 13-6
Backup Duration 13-7
Availability of Services During Backups 13-8
Interruption of a Backup Process 13-8
Scheduling Single or Recurring Backups 13-9
Starting an Immediate Backup 13-10
Checking Backup Status 13-11
Checking Log Files 13-11
Checking Scheduled Backups 13-11
Checking the Status of a Backup in Progress 13-12
Other Important Backup Tasks 13-12
Disaster Recovery on the Security Management Appliance 13-12
Restoring Other Data 13-14
Upgrading Appliance Hardware 13-14
Upgrading AsyncOS 13-16
Batch Commands for Upgrades 13-16
Determining Network Requirements 13-16
Choosing an Upgrade Method: Remote vs. Streaming 13-16
Streaming Upgrade Overview 13-17
Remote Upgrade Overview 13-17
Hardware and Software Requirements for Remote Upgrades 13-18
Hosting a Remote Upgrade Image 13-18
Important Differences in Remote Upgrading Method 13-19
Configuring Upgrade and Service Update Settings 13-19
Upgrade and Update Settings 13-20
Static Upgrade and Update Server Settings for Environments with Strict Firewall Policies 13-21
Configuring the Update and Upgrade Settings from the GUI 13-22
Before You Upgrade: Important Steps 13-24
Upgrading AsyncOS 13-24
After Upgrading 13-25
About Reverting to an Earlier Version of AsyncOS 13-25
Important Note About Reversion Impact 13-26
Reverting AsyncOS 13-26
About Updates 13-27
About URL Category Set Updates for Cisco IronPort Web Usage Controls 13-28
Configuring the Return Address for Generated Messages 13-28
Managing Alerts 13-28

Overview of Alerts 13-28
Alerts: Alert Recipients, Alert Classifications, and Severities 13-28
Alert Settings 13-29
Alert Delivery 13-29
Cisco IronPort AutoSupport 13-30
Alert Messages 13-30
Alert From Address 13-30
Alert Subject 13-30
Example Alert Message 13-30
Managing Alert Recipients 13-31
Adding New Alert Recipients 13-31
Configuring Existing Alert Recipients 13-32
Deleting Alert Recipients 13-32
Configuring Alert Settings 13-32
Editing Alert Settings 13-32
Alert Listing 13-33
Hardware Alerts 13-33
System Alerts 13-33
Changing Network Settings 13-36
Changing the System Hostname 13-36
The sethostname Command 13-36
Configuring Domain Name System Settings 13-37
Specifying DNS Servers 13-37
Multiple Entries and Priority 13-37
Using the Internet Root Servers 13-38
Reverse DNS Lookup Timeout 13-38
DNS Alert 13-38
Clearing the DNS Cache 13-38
Configuring DNS Settings via the Graphical User Interface 13-39
Configuring TCP/IP Traffic Routes 13-39
Managing Static Routes in the GUI 13-39
Modifying the Default Gateway (GUI) 13-40
Configuring the Default Gateway 13-40
Changing the Admin User's Password 13-40
Configuring the System Time 13-41
Time Zone Page 13-41
Selecting a Time Zone 13-41
Selecting a GMT Offset 13-41
Updating Time Zone Files 13-42
Editing Time Settings 13-42

Setting the System Time Using a Network Time Protocol (NTP) Server 13-42
Setting System Time Manually 13-43
Saving and Importing Configuration Settings 13-43
Managing Multiple Appliances with XML Configuration Files 13-44
Managing Configuration Files 13-44
Saving and Exporting the Current Configuration File 13-45
Loading a Configuration File 13-45
Resetting the Current Configuration 13-47
CLI Commands for Configuration Files 13-47
The showconfig, mailconfig, and saveconfig Commands 13-48
The loadconfig Command 13-49
The publishconfig Command 13-49
The backupconfig Command 13-49
Uploading Configuration Changes Using the CLI 13-49
Managing Disk Usage 13-50
Disk Space Maximums and Allocations 13-50
Reallocating Disk Space Quotas 13-51
Customizing Your View 13-52
Setting Preferences 13-52

CHAPTER 14 **Logging** 14-1
Logging Overview 14-1
Logging Versus Reporting 14-1
Log Retrieval 14-2
Filename and Directory Structure 14-2
Log Rollover and Transfer Schedule 14-2
Timestamps in Log Files 14-3
Logs Enabled by Default 14-3
Log Types 14-4
Summary of Log Types 14-4
Log Type Comparison 14-7
Using Configuration History Logs 14-7
Using CLI Audit Logs 14-8
Using FTP Server Logs 14-9
Using HTTP Logs 14-9
Using Cisco IronPort Spam Quarantine Logs 14-10
Using Cisco IronPort Spam Quarantine GUI Logs 14-10
Using Cisco IronPort Text Mail Logs 14-11
Examples of Text Mail Log Entries 14-12

Generated or Rewritten Messages 14-15
Sending a Message to the Cisco IronPort Spam Quarantine 14-15
Using NTP Logs 14-16
Using Reporting Logs 14-16
Using Reporting Query Logs 14-17
Using Safelist/Blocklist Logs 14-17
Using SMA Logs 14-18
Using Status Logs 14-19
Reading Status Logs 14-19
Using System Logs 14-21
Understanding Tracking Logs 14-21
Log Subscriptions 14-22
Configuring Log Subscriptions 14-22
Setting the Log Level 14-23
Creating a Log Subscription in the GUI 14-23
Editing Log Subscriptions 14-24
Configuring Global Settings for Logging 14-24
Logging Message Headers 14-25
Configuring Global Settings for Logging by Using the GUI 14-25
Rolling Over Log Subscriptions 14-26
Rolling Over Logs in Log Subscriptions 14-26
Rolling Over Logs Immediately Using the GUI 14-26
Rolling Over Logs Immediately via the CLI 14-26
Viewing the Most Recent Log Entries in the GUI 14-26
Viewing the Most Recent Entries in Logs (tail Command) 14-27
Example 14-27
Configuring Host Keys 14-28

CHAPTER 15 **Troubleshooting** 15-1
Accessing Customer Support from the Security Management Appliance 15-1
Technical Support 15-1
Creating a Support Request 15-1
About Remote Access for Customer Support 15-2
Enabling Remote Access for Customer Support 15-2
Packet Capture 15-3
Starting a Packet Capture 15-3
Editing Packet Capture Settings 15-4

APPENDIX A **IP Interfaces and Accessing the Appliance** A-1

IP Interfaces A-1

Configuring IP Interfaces A-2

Creating IP Interfaces Using the GUI A-3

Accessing the Appliance via FTP A-3

Secure Copy (scp) Access A-6

Accessing via a Serial Connection A-7

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

Ethernet Interfaces B-1

Selecting IP Addresses and Netmasks B-1

Sample Interface Configurations B-2

IP Addresses, Interfaces, and Routing B-3

Summary B-3

Strategies for Connecting Your Cisco IronPort Appliance B-3

APPENDIX C **Firewall Information** C-1

APPENDIX D **Examples** D-1

Web Security Appliance Examples D-1

Example 1: Investigating a User D-1

Related Topics D-5

Example 2: Tracking a URL D-5

Related Topics D-6

Example 3: Investigating Top URL Categories Visited D-6

Related Topics D-8

APPENDIX E **End User License Agreement** E-1

Cisco Systems End User License Agreement E-1

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

INDEX

CHAPTER 1

Getting Started

This guide provides instructions for setting up, administering, and monitoring the Cisco IronPort Security Management appliance. These instructions are designed for an experienced system administrator with knowledge of networking and email and web administration.

- What's New in This Release
- Where to Find More Information
- Cisco Welcomes Your Comments
- Security Management Appliance Overview

What's New in This Release

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

http://www.cisco.com/en/US/products/ps10155/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.

Feature	Description
New Features:	
Reporting and Tracking support for new web security features	A new report has been added for SOCKS Proxy transactions, including destinations and users. A new tab for SOCKS Proxy transactions has been added to the Web Tracking page. For information, see SOCKS Proxy Report, page 5-46 and Searching for Transactions Processed by the SOCKS Proxy, page 5-55. Also, transactions blocked because of OCSP policies are included in existing reports and web tracking data.
Configuration Master supports new features	A new Configuration Master supports the new SOCKS Proxy feature in AsyncOS for Web Security 7.7.
Support for Web Security appliances in FIPS mode	This release can manage supported Web Security appliances that are running in Federal Information Processing Standard (FIPS) mode.

Feature	Description
Enhancements:	
Supported browsers	You can now use Internet Explorer 9 to access the appliance via the Web.

Where to Find More Information

Cisco IronPort offers the following resources to learn more about the Security Management appliance and related products:

- Documentation, page 1-2
- Training, page 1-3
- Knowledge Base, page 1-3
- Cisco Support Community, page 1-4
- Customer Support, page 1-4
- Third Party Contributors, page 1-3

Documentation

Documentation for this product and related products is available at the following locations:

Documentation For Cisco IronPort Products:	Is Located At:
Security Management appliances	http://www.cisco.com/en/US/products/ps10155/tsd_products_support_series_home.html
Email Security appliances and the CLI reference guide	http://www.cisco.com/en/US/products/ps10154/tsd_products_support_series_home.html
Web Security appliances	http://www.cisco.com/en/US/products/ps10164/tsd_products_support_series_home.html
Cisco IronPort Encryption	http://www.cisco.com/en/US/partner/products/ps10602/tsd_products_support_series_home.html

The guide is distributed as PDF and HTML files. The electronic versions of the guide are available on the Cisco IronPort Customer Support site. 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.

The documentation set for Cisco IronPort appliances includes the following documents and books (not all types are available for all appliances and releases):

- Release Notes for all products
- The *Quick Start Guide* for the Security Management appliance

- *Cisco IronPort AsyncOS for Security Management User Guide* (this book)
- *Cisco IronPort AsyncOS for Web Security User Guide*
- Cisco IronPort AsyncOS for Email Security user guides:
  - *Cisco IronPort AsyncOS for Email Security Configuration Guide*
  - *Cisco IronPort AsyncOS for Email Security Advanced Configuration Guide*
  - *Cisco IronPort AsyncOS for Email Security Daily Management Guide*
- *Cisco IronPort AsyncOS CLI Reference Guide*

Third Party Contributors

Some software included within Cisco IronPort 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 IronPort license agreements.

Information about third-party licenses is available in a Licensing document at: http://www.cisco.com/en/US/products/ps10155/prod_release_notes_list.html and at https://support.ironport.com/3rdparty/AsyncOS_User_Guide-1-1.html.

Portions of the software within Cisco IronPort 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/web/learning/le31/email_sec/index.html
- http://www.cisco.com/web/learning/training-index.html
- stbu-trg@cisco.com

Knowledge Base

You can access the Cisco IronPort Knowledge Base on the Cisco IronPort Customer Support site at the following URL:

`http://www.cisco.com/web/ironport/knowledgebase.html`

Note You need a Cisco.com User ID to access the site. If you do not have a Cisco.com User ID, you can register for one here: `https://tools.cisco.com/RPF/register/register.do`

The Knowledge Base contains a wealth of information on topics related to Cisco IronPort products.

Articles generally fall into one of the following categories:

- **How-To.** These articles explain how to do something with a Cisco IronPort product. For example, a how-to article might explain the procedures for backing up and restoring a database for an appliance.

- **Problem-and-Solution.** A problem-and-solution article addresses a particular error or issue that you might encounter when using a Cisco IronPort product. For example, a problem-and-solution article might explain what to do if a specific error message is displayed when you upgrade to a new version of the product.
- **Reference.** Reference articles typically provide lists of information, such as the error codes associated with a particular piece of hardware.
- **Troubleshooting.** Troubleshooting articles explain how to analyze and resolve common issues related to Cisco IronPort products. For example, a troubleshooting article might provide steps to follow if you are having problems with DNS.

Cisco Support Community

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

Access the Cisco Support Community at the following URLs:

- For email security and associated management:

  https://supportforums.cisco.com/community/netpro/security/email

- For web security and associated management:

  https://supportforums.cisco.com/community/netpro/security/web

Customer Support

Use the following methods to obtain support:

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

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

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

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

See also Accessing Customer Support from the Security Management Appliance, page 15-1.

Cisco Welcomes Your Comments

The Cisco IronPort 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.

Security Management Appliance Overview

AsyncOS for 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 reporting:** Run reports on aggregated data from multiple Email and Web Security appliances. The same reporting features available on individual appliances are available on Security Management appliances. In addition, there are several extended reports for web security that are uniquely available on the Security Management appliance.
- **Centralized tracking:** Use a single interface to track email messages and web transactions that were processed by multiple Email and Web Security appliances.
- **Centralized Configuration Management:** For simplicity and consistency, manage policy definition and policy deployment for up to 150 Web Security appliances. Policies can be pushed from the Security Management appliance to appliances running multiple AsyncOS versions.
- **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.

Note The Security Management appliance is not involved in centralized email management, or ‘clustering’ of Email Security appliances.

CHAPTER 2

Setup, Installation, and Basic Configuration

- Solution Deployment Overview, page 2-1
- SMA Compatibility Matrix, page 2-2
- Installation Planning, page 2-2
- Preparing for Setup, page 2-4
- Accessing the Security Management Appliance, page 2-6
- Running the System Setup Wizard, page 2-8
- About Adding Managed Appliances, page 2-12
- Configuring Services on the Security Management Appliance, page 2-14
- Committing and Abandoning Configuration Changes, page 2-14

Solution Deployment Overview

To configure your centrally-managed Cisco IronPort 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.	See the SMA Compatibility Matrix, page 2-2.
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 *Cisco IronPort AsyncOS for Email Security* documentation.

	On These Appliances	Do This	More Information
Step 3	**Web Security appliances**	Before you introduce centralized services to your environment, configure at least one Web Security appliance to provide the security features you want, and verify that all features are working as expected.	See the *Cisco IronPort AsyncOS for Web Security User Guide.*
Step 4	**Security Management appliance**	Set up the appliance and run the System Setup Wizard.	See the "Installation Planning" section on page 2-2, "Preparing for Setup" section on page 2-4 and the "Running the System Setup Wizard" section on page 2-8.
Step 5	**All appliances**	Configure each centralized service that you want to deploy.	Start with the "Configuring Services on the Security Management Appliance" section on page 2-14.

SMA Compatibility Matrix

For compatibility of your Security Management appliance with Email Security appliances and Web Security appliances, and for compatibility of configuration files when importing and publishing Web Security appliance configurations, see the Release Notes for your release at http://www.cisco.com/en/US/products/ps10155/prod_release_notes_list.html.

Installation Planning

- Network Planning, page 2-2
- About Integrating a Security Management Appliance with Email Security Appliances, page 2-3
- Deployments with Centralized Management of Email Security Appliances, page 2-3

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 (see Figure 2-1).

Figure 2-1 Typical Network Configuration Incorporating the Security Management appliance

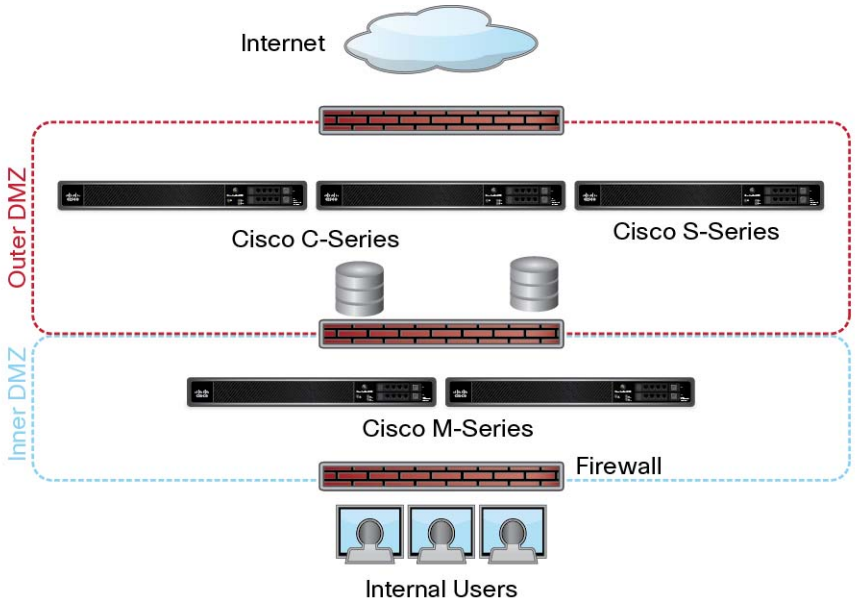

Figure 2-1 shows a typical network configuration incorporating the Security Management appliance and multiple DMZs. You deploy the Security Management appliance outside your DMZ, in your internal networks. All connections are initiated by the Security Management appliances (Cisco IronPort M-Series) to the managed Email Security appliances (Cisco IronPort C-Series) and managed Web Security appliances (Cisco IronPort S-Series).

Corporate data centers can share a Security Management appliance to perform centralized reporting and message tracking for multiple Web and Email Security appliances, and centralized policy configuration for multiple Web Security appliances. The Security Management appliance can also be used as an external Cisco IronPort Spam Quarantine.

After you connect the Email Security appliance and the Web Security appliances 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 and web usage can be determined.

About Integrating a Security Management Appliance with Email Security Appliances

For an excellent discussion on integrating the Security Management appliance with your Email Security appliances, see the "Cisco IronPort M-Series Security Management Appliance" chapter in the *Cisco IronPort AsyncOS for Email Security Configuration Guide.* ("M-Series appliance" is another term for the Security Management appliance.)

Deployments with Centralized Management of Email Security Appliances

The Security Management appliance cannot be placed in a cluster. However, clustered Email Security appliances can deliver messages to the Security Management appliance for centralized reporting and tracking and to store the messages in an external spam quarantine.

Preparing for Setup

Before you run the System Setup Wizard:

Step 1 Review the latest release notes for your product. See Documentation, page 1-2.

Step 2 Verify that the components of your security solution are compatible. See the SMA Compatibility Matrix, page 2-2.

Step 3 Ensure that your network and physical space are ready to support this deployment. See Installation Planning, page 2-2.

Step 4 Physically set up and connect the Security Management appliance. See Physically Setting Up and Connecting the Appliance, page 2-4.

Step 5 Determine network and IP address assignments. See Determining Network and IP Address Assignments, page 2-4.

Step 6 Gather information about your system setup. See Gathering the Setup Information, page 2-5.

Physically Setting Up and Connecting the Appliance

Before you follow the procedures in this chapter, complete the steps described in the *Cisco IronPort M-Series Quickstart 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 Cisco IronPort 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 Appendix B, "Assigning Network and IP Addresses."

Note If you are running a firewall on your network between the Internet and the Cisco IronPort appliance, it may be necessary to open specific ports for the Cisco IronPort appliance to work properly. For more information on firewalls, see Appendix C, "Firewall Information."

Note 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 Cisco IronPort AsyncOS for Email Security.

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 Appendix B, "Assigning Network and IP Addresses," for detailed information about network and IP addresses.

Table 2-1 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 Cisco IronPort 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 2-2 Supported Browsers and Releases

Browser	Windows XP	Windows 7	MacOS X
Safari	—	—	4.0 and later
Google Chrome	Latest stable release	Latest stable release	Latest stable release
Microsoft Internet Explorer	7.0, 8.0	8.0, 9.0	—
Mozilla Firefox	10.x and later	10.x and later	10.x and later

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

Accessing the Web Interface

Procedure

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**

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 Cisco IronPort Spam Quarantine end user interface, available by default on port 82. The Cisco IronPort 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 Cisco IronPort 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.

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. To change the timeout limit, see Configuring the Web UI Session Timeout, page 12-20.

Accessing the Security Management Appliance 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 IronPort 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 *Cisco IronPort AsyncOS CLI Reference Guide*.

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, page 13-52.

- 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" section on page 2-4.

Warning **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.

Warning **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 the session times out while you are running the System Setup Wizard, you need to start over from the beginning.

To change the session timeout limit, see Configuring the Web UI Session Timeout, page 12-20.

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, page 12-20.

Overview of the System Setup Wizard

Procedure

Step 1 Reviewing the end user license agreement

Step 2 Configuring the following system settings:

- Notification settings and AutoSupport

- System time settings
- Admin password

Step 3 Configuring the following network settings:

- Hostname of the appliance
- IP address, network mask, and gateway of the appliance
- Default router and DNS settings

Step 4 Reviewing your configuration

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.

Launch the System Setup Wizard

To launch the wizard, log in to the GUI as described in the "Accessing the Web Interface" section on page 2-6. 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

Figure 2-2 Configuring System Settings with the System Configuration Page

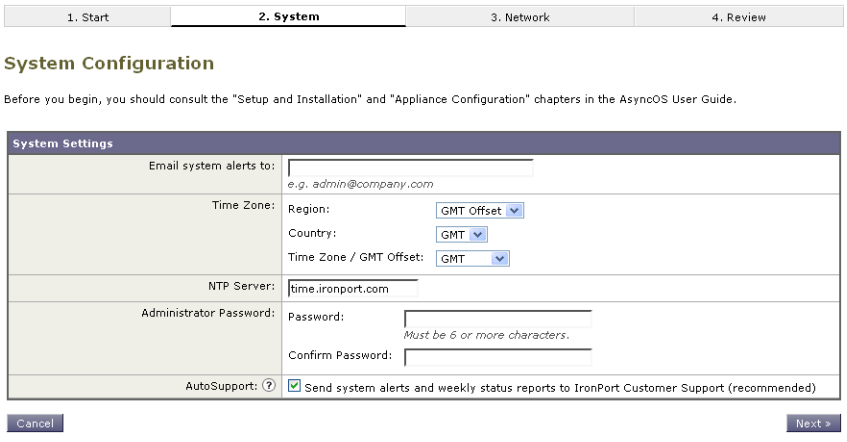

Entering an Email Address for System Alerts

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" section on page 13-28.

Setting the Time

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

You can set the system clock time manually, or you can use Network Time Protocol (NTP) to synchronize time with other servers on your network or the Internet. By default, the Cisco IronPort Systems time servers (time.ironport.com) are added as an entry to synchronize the time on your Security Management 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" section on page 13-41.

Note When gathering data for reports, the Security Management appliance applies a time stamp on the data. The time stamp is applied using the configuration settings that you implemented from the steps in the "Configuring the System Time" section on page 13-41.
For more information on how the Security Management appliance gathers data, see the "How the Security Appliance Gathers Data for Reports" section on page 3-2.

Setting the Password

You must change the password for the AsyncOS admin account. The new password must be six characters or longer. 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 Cisco IronPort AutoSupport feature (enabled by default) notifies Cisco IronPort Customer Support about issues with the Security Management appliance so that they can provide optimal support. For more information, see the "Cisco IronPort AutoSupport" section on page 13-30.

Configure the Network Settings

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

Figure 2-3 Configuring Network Settings

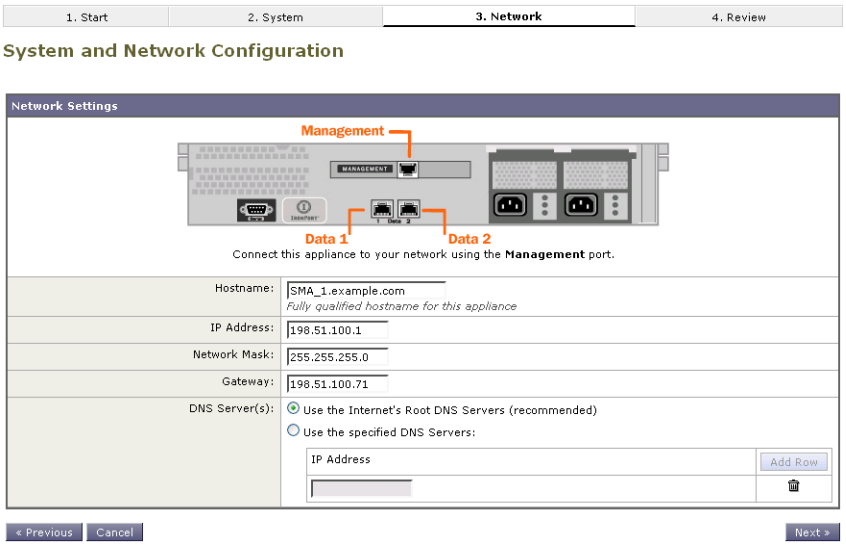

Note Verify that you have connected the Security Management appliance to your network through the Management port.

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 router (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" section on page 13-37.

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.

Figure 2-4 Reviewing Your Configuration

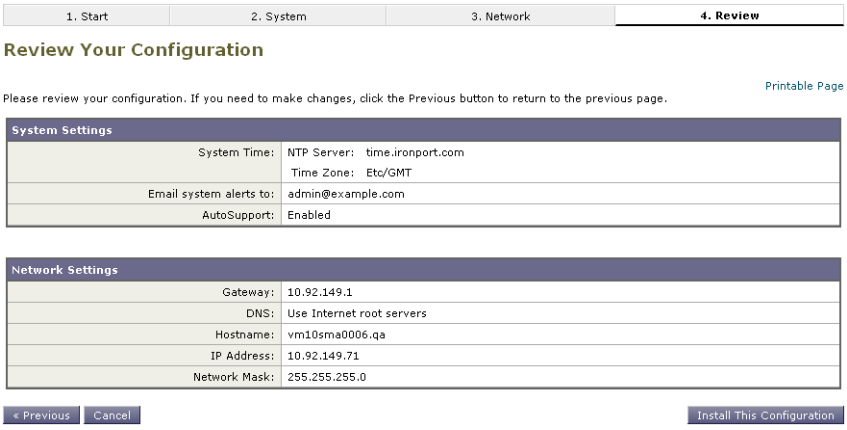

After you have reviewed the information, click **Install This Configuration**. Then click **Install** in the confirmation dialog box that appears.

Proceeding to the Next Steps

If the System Setup Wizard properly installs the configuration on the Security Management appliance, the **System Setup Next Steps** page appears.

Figure 2-5 System Setup —Next Steps

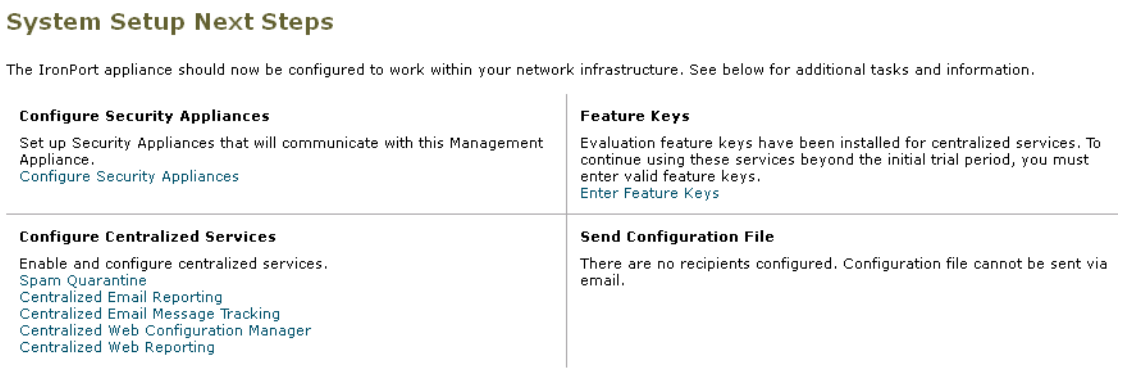

Click on any of the links on the System Setup Next Steps page to proceed with the configuration of your Cisco IronPort appliances.

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.

To simplify configuration and troubleshooting, we recommend that you follow the process outlined in Solution Deployment Overview, page 2-1.

About Adding Managed Appliances

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

Supported Email and Web Security appliances are shown in the SMA Compatibility Matrix, page 2-2.

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. For example, you add an appliance using the Add Web Security appliance page, the Security Management appliance checks the product name of the remote appliance to make sure that it is a Web Security appliance and not an Email Security appliance. 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, page 4-3
- Adding the Centralized Message Tracking Service to Each Managed Email Security Appliance, page 6-3
- Adding the Centralized Spam Quarantine Service to Each Managed Email Security Appliance, page 7-6
- Adding the Centralized Web Reporting Service to Each Managed Web Security Appliance, page 5-3
- Adding Web Security Appliances and Associating Them with Configuration Master Versions, page 8-5

Editing Managed Appliance Configurations

Procedure

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

Step 2 In the Security Appliance section, click on the name of the appliance you want to edit.

Step 3 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 a Web Security appliance, the publish history for the appliance will be lost, and publishing errors will occur if the Web Security appliance is currently selected for a scheduled publish job. (This does not affect scheduled publish jobs that are set to use all assigned appliances.) If you change the IP address of an Email Security appliance, the tracking availability data for the appliance will be lost.

Step 4 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

Procedure

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

Step 2 In the Security Appliances section, and click the trash can icon in the row for the managed appliance that you want to delete.

Step 3 In the confirmation dialog box, click **Delete**.

Step 4 Submit and commit your changes.

Configuring Services on the Security Management Appliance

To configure and use one or more security services on the Security Management appliance, see the following information:

Email security services:

- Chapter 4, "Using Centralized Email Security Reporting"
- Chapter 6, "Tracking Email Messages"
- Chapter 7, "Managing the Cisco IronPort Spam Quarantine"

Web security services:

- Chapter 5, "Using Centralized Web Reporting and Tracking"
- Chapter 8, "Managing Web Security Appliances"

Committing and Abandoning Configuration Changes

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

Figure 2-6 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**.

CHAPTER 3

Working With Reports

Information in this chapter applies to both email and web reports unless otherwise noted:

- Ways to View Reporting Data, page 3-1
- How the Security Appliance Gathers Data for Reports, page 3-2
- Customizing Your View of Report Data, page 3-3
- Viewing Details of Transactions Included in Reports, page 3-6
- Improving Performance of Email Reports, page 3-7
- Printing and Exporting Reporting and Tracking Data, page 3-8
- Subdomains vs. Second-Level Domains in Reporting and Tracking, page 3-10
- Email and Web Reports, page 3-10

Ways to View Reporting Data

Table 3-1 Ways To View Reporting Data

To	See
View and customize web-based interactive report pages	• Customizing Your View of Report Data, page 3-3 • Chapter 4, "Using Centralized Email Security Reporting" • Chapter 5, "Using Centralized Web Reporting and Tracking"
Automatically generate recurring PDF or CSV reports	• Scheduling Email Reports, page 4-55 • Scheduling Web Reports, page 5-61
Generate a PDF or CSV report on demand	• Generating Email Reports On Demand, page 4-57 • Generating Web Reports on Demand, page 5-65.

Table 3-1 Ways To View Reporting Data

To	See
Export raw data as a CSV (Comma-separated values) file	• Printing and Exporting Reporting and Tracking Data, page 3-8 • Exporting Report Data as a Comma-Separated Values (CSV) File, page 3-9
Generate a PDF of report data	Printing and Exporting Reporting and Tracking Data, page 3-8
Email report information to yourself and other people	• Generating Email Reports On Demand, page 4-57 • Scheduling Email Reports, page 4-55 • Generating Web Reports on Demand, page 5-65. • Scheduling Web Reports, page 5-61
View archived copies of scheduled and on-demand reports until they are purged from the system	Viewing and Managing Archived Web Reports, page 5-66
Find information about specific transactions	• Viewing Details of Transactions Included in Reports, page 3-6

Note For differences between logging and reporting, see Logging Versus Reporting, page 14-1.

How the Security 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.

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 on setting the time on your Security Management appliance, see the "Configuring the System Time" section on page 13-41.

How Reporting Data is Stored

All of the appliances store reporting data. Table 3-2 shows what time periods that each appliance stores data.

Table 3-2 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	See Viewing Reporting Data for an Appliance or Reporting Group, page 3-4
Specify a time range	See Choosing a Time Range for Reports, page 3-4
(For Web reports) Choose which data to chart	See (Web Reports Only) Choosing Which Data to Chart, page 3-5
Customize tables	See Customizing Tables on Report Pages, page 3-6
Search for specific information or a subset of data to view	• For Email reports, see Searching and the Interactive Email Report Pages, page 4-5. • 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, page 13-52

Note Not all customization features are available for every report.

Viewing Reporting Data for an Appliance or Reporting Group

For Overview reports for Email and Web, and for 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, page 4-4, 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.

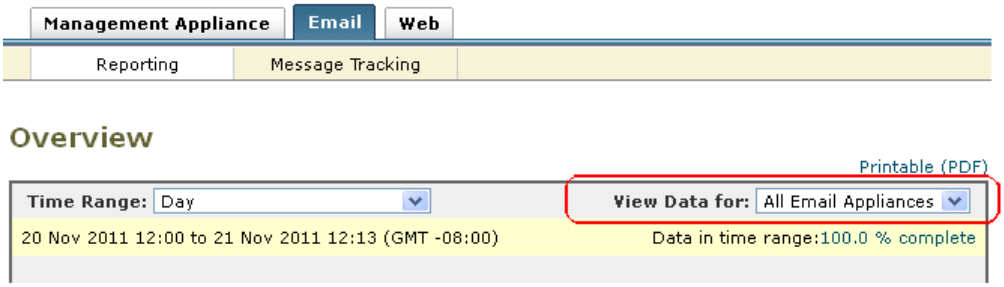

If you are viewing report data on a Security Management appliance to which you have recently backed up data from another Security Management appliance, you must first add (but do not establish a connection to) each appliance in Management Appliance > Centralized Services > Security Appliances.

Choosing a Time Range for Reports

Most interactive report pages allow you to choose a Time Range for the data to include. The time range that you select is used throughout 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 and Web reporting on the Security Management appliance:

Table 3-3 Time Range Options for Reports

Option	Description	SMA Email reports	ESA	SMA Web reports	WSA
Hour	The previous 60 minutes, plus up to 5 additional minutes		•		•
Day	The previous 24 hours	•	•	•	•
Week	The previous seven days, including the elapsed hours of the current day	•	•	•	•
30 days	The previous 30 days, including the elapsed hours of the current day	•	•	•	•
90 days	The previous 90 days, including the elapsed hours of the current day	•	•	•	
Year	The last 12 months plus the elapsed days of the current month	•			

Table 3-3 Time Range Options for Reports

Option	Description	SMA Email reports	ESA	SMA Web reports	WSA
Yesterday	24 hours (00:00 to 23:59) of the previous day, using the time zone defined on the appliance	•	•	•	•
Previous Calendar Month	00:00 of the first day of the month to 23:59 of the last day of the month	•	•	•	
Custom Range	The time range that you specify. Select this option to choose start and end dates and times.	•	•	•	•

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, page 13-52.

(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. For explanations of the column headings, see Table Column Descriptions for Web Reports, page 5-10.

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

Procedure

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 3-4 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. For column descriptions, see: • Table Column Descriptions for Email Reporting Pages, page 4-8 • Table Column Descriptions for Web Reports, page 5-10.
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.	—
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, page 13-52.
View details about a table entry, where available	Click a blue entry in the table	See also Viewing Details of Transactions Included in Reports, page 3-6.
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 Understanding the Web Reporting Pages, page 5-7.

Viewing Details of Transactions Included in Reports

Procedure

Step 1 Click any blue number in a table on a report page.

(Not all tables have these links.)

The transactions included in that number are displayed in Web Tracking.

Step 2 Scroll down to see the list of transactions.

Related Topics

- Web Tracking, page 5-50

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 and Tracking Data

Table 3-5 Printing and Exporting Report Data

To Get This	PDF	CSV	Do This	Notes
A PDF of an interactive report page	•		Click the **Printable (PDF)** link at the top-right of an interactive report page.	The PDF reflects the customizations that you are currently viewing. PDFs are formatted to be printer-friendly.
A PDF of report data	•		Create a scheduled or on-demand report. See: • Generating Email Reports On Demand, page 4-57 • Scheduling Email Reports, page 4-55 • Generating Web Reports on Demand, page 5-65. • Scheduling Web Reports, page 5-61	—
Raw data See also Exporting Report Data as a Comma-Separated Values (CSV) File, page 3-9.		•	Click the **Export** link below the chart or table.	The CSV file contains all applicable data up to the limit, not just the data visible in the chart or table.
		•	Create a scheduled or on-demand report. See: • Generating Email Reports On Demand, page 4-57 • Scheduling Email Reports, page 4-55 • Generating Web Reports on Demand, page 5-65. • Scheduling Web Reports, page 5-61	Each CSV file may contain up to 100 rows. If a report contains more than one table, a separate CSV file will be created for each table. Some extended reports are not available in CSV format.
Reports in different languages	•		Choose the desired Report Language when you schedule a report or create one on demand.	To generate PDFs in Chinese, Japanese, or Korean on Windows computers, you must also download the applicable Font Pack from Adobe.com and install it on your local computer.

Table 3-5 Printing and Exporting Report Data (continued)

To Get This	PDF	CSV	Do This	Notes
(Web Security) A custom subset of report data, for example data for a particular user.	•	•	In Web Tracking, perform a search, then click the Printable Download link on the Web Tracking page. Choose PDF or CSV format.	PDFs may not include all information available on the web page. Specifically, PDFs include: • Up to 1,000 transactions. • If you display details, up to 100 related transactions. • Up to 3000 characters per related transaction. CSV files include all raw data matching the search criteria.
(Email Security) A custom subset of data, for example data for a particular user.		•	In Message Tracking, perform your search, then click the Export link above the search results.	Message tracking search results and CSV file include up to 1000 results, depending on the limit you specified in your search criteria.

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 and Tracking Data, page 3-8.

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 3-6 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.

Table 3-6 Viewing Raw Data Entries (continued)

Category Header	Value	Description
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`.

Email and Web Reports

For information specific to Email reports, see Chapter 4, “Using Centralized Email Security Reporting.”

For information specific to Web reports, see Chapter 5, “Using Centralized Web Reporting and Tracking.”

CHAPTER 4

Using Centralized Email Security Reporting

- Centralized Email Reporting Overview, page 4-1
- Setting Up Centralized Email Reporting, page 4-2
- Working with Email Report Data, page 4-4
- Understanding the Email Reporting Pages, page 4-6
- About Scheduled and On-Demand Email Reports, page 4-50
- Generating Email Reports On Demand, page 4-57
- Scheduling Email Reports, page 4-55
- Viewing and Managing Archived Email Reports, page 4-58

Centralized Email Reporting Overview

The Email reporting functionality aggregates 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 *Cisco IronPort AsyncOS for Email Security Configuration Guide.*

Setting Up Centralized Email Reporting

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

- Enabling Centralized Email Reporting on the Security Management Appliance, page 4-2.
- Adding the Centralized Email Reporting Service to Each Managed Email Security Appliance, page 4-3.
- Creating Email Reporting Groups, page 4-4
- Enabling Centralized Email Reporting on Email Security Appliances, page 4-4

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 Usage" section on page 13-50.

Procedure

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 Usage" section on page 13-50, 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.

Procedure

Step 1 On the Security Management appliance, 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 Cisco IronPort 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, page 4-3.

Procedure

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

Step 2 Click **Add Group**.

Step 3 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 4 Click **Add** to add the appliances to the Group Members list.

Step 5 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 *Cisco IronPort AsyncOS for Email Security Configuration Guide*.

Working with Email Report Data

- For options for accessing and viewing report data, see "Ways to View Reporting Data" section on page 3-1.
- To customize your view of report data, see Customizing Your View of Report Data, page 3-3
- To search for specific information within your data, see Searching and the Interactive Email Report Pages, page 4-5.

- To print or export report information, see Printing and Exporting Reporting and Tracking Data, page 3-8
- To understand the various interactive report pages, see Understanding the Email Reporting Pages, page 4-6.
- To generate a report on demand, see Generating Email Reports On Demand, page 4-57.
- To schedule reports to run automatically at intervals and times that you specify, see Scheduling Email Reports, page 4-55.
- To view archived on-demand and scheduled reports, see Viewing and Managing Archived Email Reports, page 4-58.
- For background information, How the Security Appliance Gathers Data for Reports, page 3-2.
- To improve performance when working with large amounts of data, see "Improving Performance of Email Reports" section on page 3-7.
- To get details about an entity or number that appears as a blue link in a chart or table, click the entity or number.

Searching and the Interactive Email Report Pages

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

The following picture shows the **Search For** drop-down menu.

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

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

Table 4-1 Email Reporting Tab Options

Email Reporting Menu	Action
Email Reporting Overview Page	The Overview page provides a synopsis of the activity on your Cisco IronPort Email appliances. It includes graphs and summary tables for the incoming and outgoing messages. For more information, see the "Email Reporting Overview Page" section on page 4-10.
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" section on page 4-14.
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" section on page 4-24.
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" section on page 4-26.
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" section on page 4-27.
DLP Incident Summary Page	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 Incident Summary Page" section on page 4-30.

Table 4-1 Email Reporting Tab Options

Email Reporting Menu	Action
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" section on page 4-32.
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" section on page 4-34.
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" section on page 4-36.
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" section on page 4-39.
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 messages per sender. For more information, see the "Rate Limits Page" section on page 4-40.
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" section on page 4-41.
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" section on page 4-43.

Table 4-1 Email Reporting Tab Options

Email Reporting Menu	Action
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" section on page 4-49.
Scheduling Email Reports	Allows you to schedule reports for a specified time range. For more information, see the "Scheduling Email Reports" section on page 4-55.
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" section on page 4-58. Also allows you to generate on-demand reports. See "Generating Email Reports On Demand" section on page 4-57.

Table Column Descriptions for Email Reporting Pages

Table 4-2 Table Column Descriptions for Email Reporting Pages

Column Name	Description
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.
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.

Table 4-2 Table Column Descriptions for Email Reporting Pages

Column Name	Description
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.
Clean	All clean messages.
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.

Table 4-2 Table Column Descriptions for Email Reporting Pages

Column Name	Description
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.
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.

Figure 4-1 The Email > Reporting > Overview Page

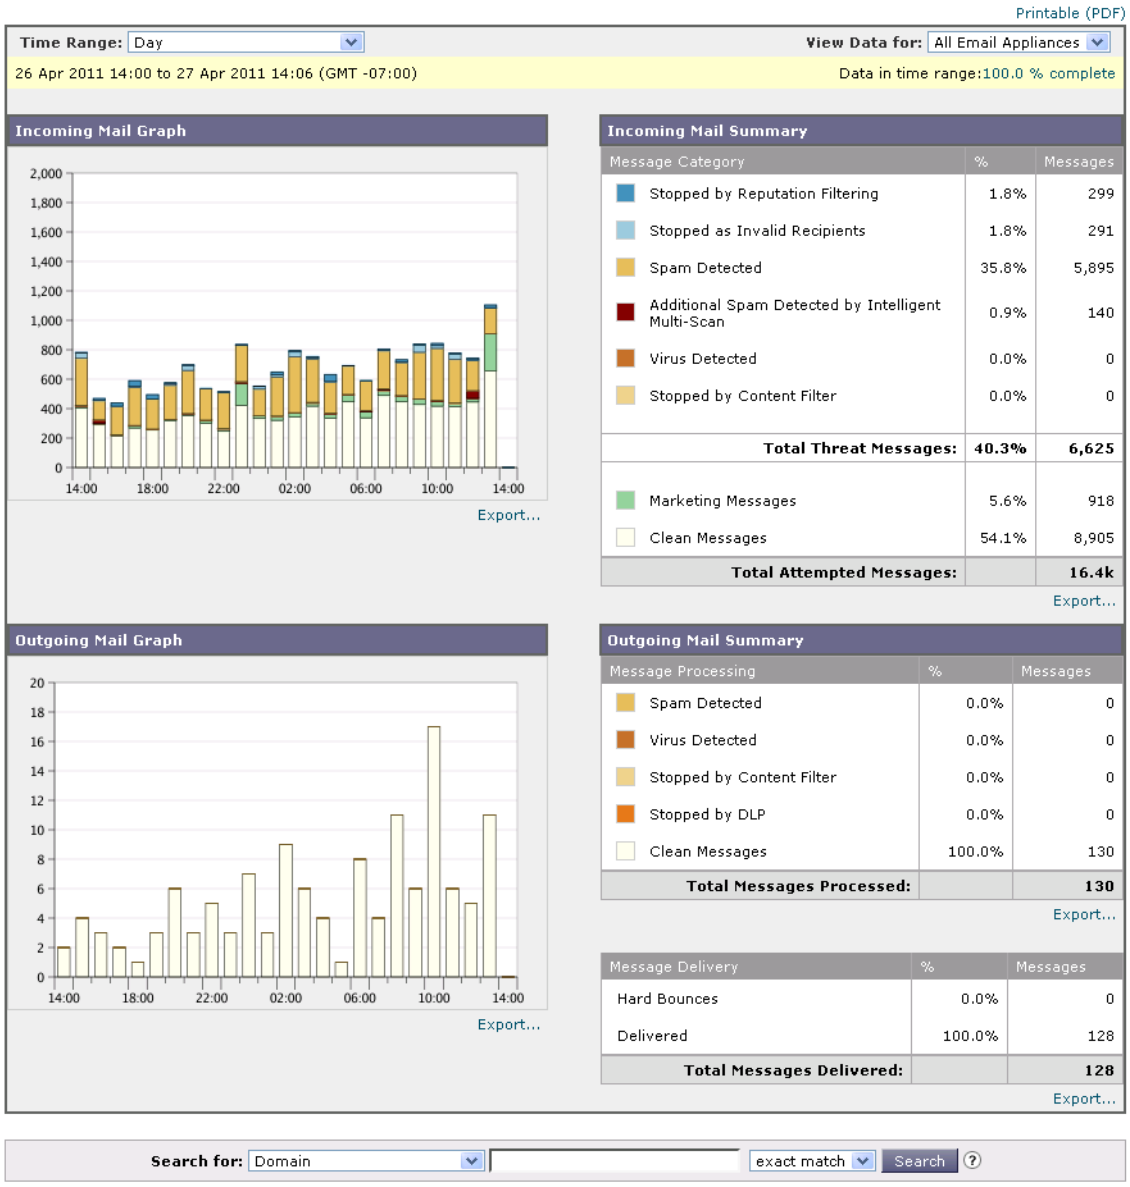

At a high level the **Overview** page shows you the incoming and outgoing mail graphs, and well as incoming and outgoing mail summaries.

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 It is important to remember that the Domain-Based Executive Summary Report, and the Executive Summary report are based on the Email Reporting Overview Page. The Domain-Based Executive Summary report is limited to a group of specified domains. For information scheduling reports, see the "Scheduling Email Reports" section on page 4-55.

Table 4-3 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" section on page 3-4.
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, page 3-4.
Incoming Mail Graph	The Incoming Mail Graph displays a visual graph of the breakdown of incoming mail in real time.
Outgoing Mail Graph	The Outgoing Mail graph displays a visual graph of the breakdown of outgoing mail on the appliance.
Incoming Mail Summary	The Incoming Mail Summary shows the percentages and the number of messages that were stopped by reputation filtering (SBRS), stopped as invalid recipient, spam detected, virus detected, and stopped by content filter, and those considered "clean."
Outgoing Mail Summary	The Outgoing Mail Summary section includes information about the outgoing threat and clean messages. It also includes a breakdown of the delivered versus hard-bounced messages.

How Incoming Mail Messages are Counted

AsyncOS counts incoming mail 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 reputation filtering do not actually enter the work queue, the appliance does not have access to the list of recipients for an incoming message. In this case, a multiplier is used to estimate the number of recipients. This multiplier was determined by Cisco IronPort Systems and 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 various verdicts follow these rules of precedence:

- Outbreak Filters quarantining
  (in this case the message is not counted until it is released from the quarantine and again processed through the work queue)
- 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.

Categorizing Email Messages on Overview Page

Messages reported on the Overview page are categorized as follows:

Table 4-4 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" section on page 4-12) plus all recipients blocked by recipient throttling. 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.
Virus Messages Detected	The total count and percentage of messages detected as virus positive and not also spam.
Stopped by Content Filter	The total count of messages that were stopped by a content filter.
Marketing Messages	The total count and percentage of messages detected as unwanted marketing messages. This list item appears on the page only if marketing data are present in the system.
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.

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.

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** page consists of two main sections: the mail trend graphs summarizing the top senders (by total threat messages and by total clean messages) and the Incoming Mail Details interactive table.

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, page 4-5.
- 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" section on page 4-23 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 (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 Cisco IronPort 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 Cisco IronPort 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 Mai**l 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" section on page 4-19.

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 Groups Report Page, page 4-23.

Categorizing Email Messages on Incoming Mail Page

Messages reported on the Incoming Mail page are categorized as follows:

Table 4-5 Email Categories on Incoming Mail 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" section on page 4-12) 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.
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.

Table 4-5 Email Categories on Incoming Mail Page

Category	Description
Virus Messages Detected	The total count and percentage of messages detected as virus positive and not also spam.
Stopped by Content Filter	The total count of messages that were stopped by a content filter.
Marketing Messages	The total count and percentage of messages detected as unwanted marketing messages. This list item appears on the page only if marketing data are present in the system
Clean Messages Accepted	Mail that is accepted and that is deemed to be virus and spam free — the most accurate representation of clean messages accepted when taking per recipient scanning actions (such as splintered messages being processed by separate mail policies) into account. However, because messages that are marked as spam or virus positive and still delivered are not counted, the actual number of messages delivered may differ from the clean message count.

Note 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.

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, page 4-6.

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

Figure 4-2 Incoming Mail Page: IP Address View

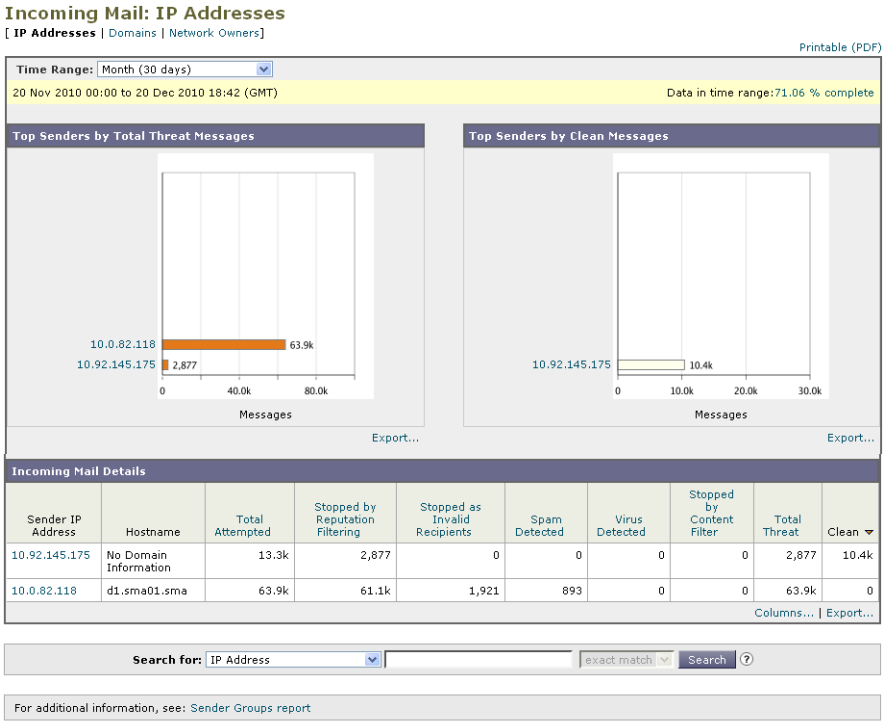

See the "Incoming Mail Details Table" section on page 4-19 for an explanation of the data included in the Incoming Mail Details interactive table.

In this example, the **Domai**n view is selected.

Figure 4-3 Incoming Mail Page: Domain View

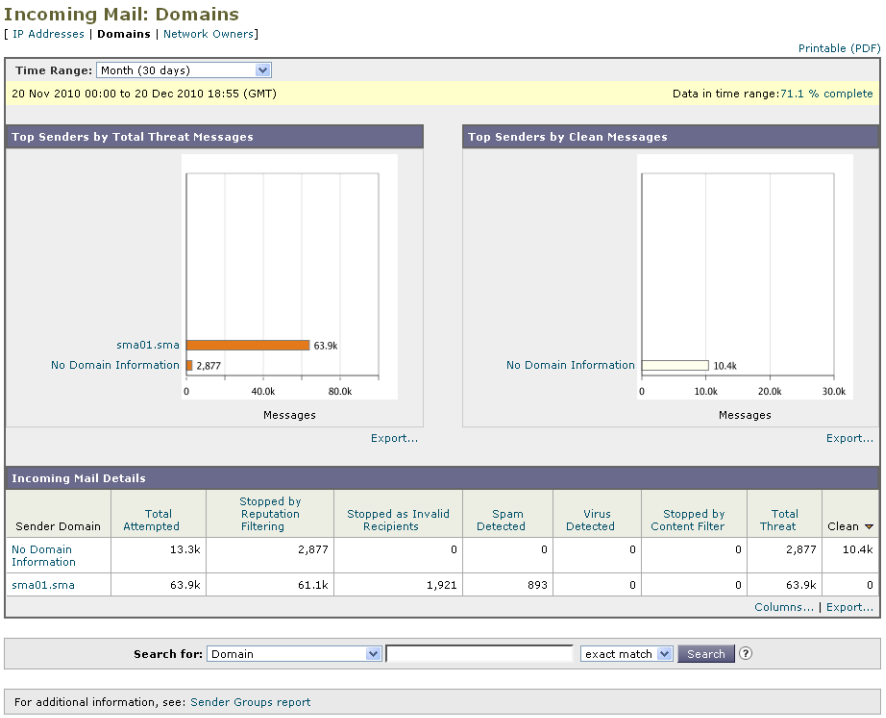

From the **Incoming Mail** page you can also generate a PDF or export raw data to a CSV file. For information on printing or exporting a file, see the "Understanding the Email Reporting Pages" section on page 4-6.

Note You can generate a scheduled report for the Incoming Mail report page. See the "Scheduling Email Reports" section on page 4-55.

"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 *Cisco IronPort AsyncOS for Email Security Configuration Guide*.

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" section on page 3-4.

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 *Cisco IronPort AsyncOS for Email Security Configuration Guide*.

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 Cisco IronPort 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" section on page 4-19.

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" section on page 4-23.

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.

Figure 4-4 Domains Listing for Network Owner

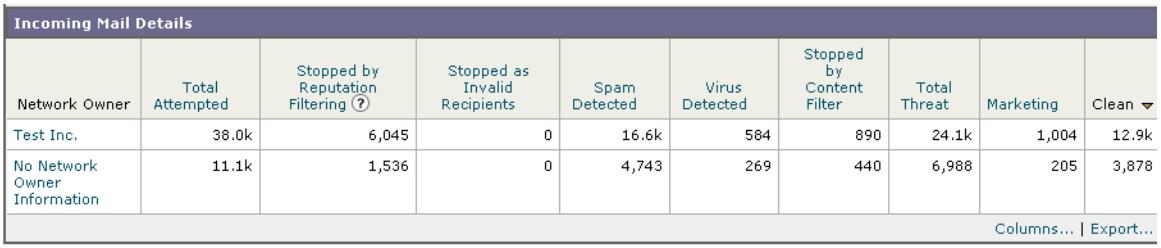
Incoming Mail Details

Network Owner	Total Attempted	Stopped by Reputation Filtering	Stopped as Invalid Recipients	Spam Detected	Virus Detected	Stopped by Content Filter	Total Threat	Marketing	Clean
Test Inc.	38.0k	6,045	0	16.6k	584	890	24.1k	1,004	12.9k
No Network Owner Information	11.1k	1,536	0	4,743	269	440	6,988	205	3,878

Columns... | Export...

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

- The global information from the Cisco IronPort 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 Cisco IronPort 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.

Figure 4-5 Current Information for Network Owner

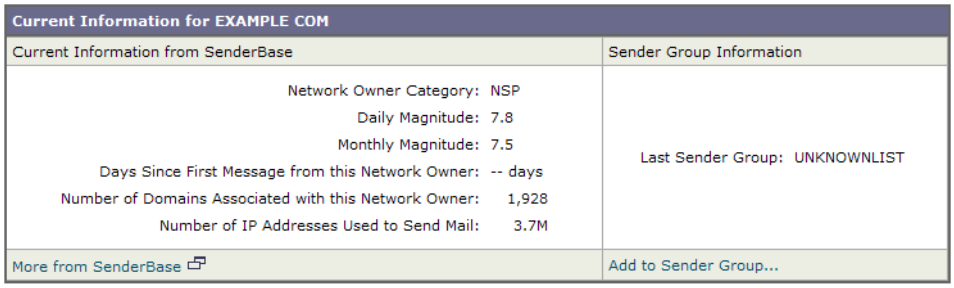

Figure 4-6 Domain Profile Page

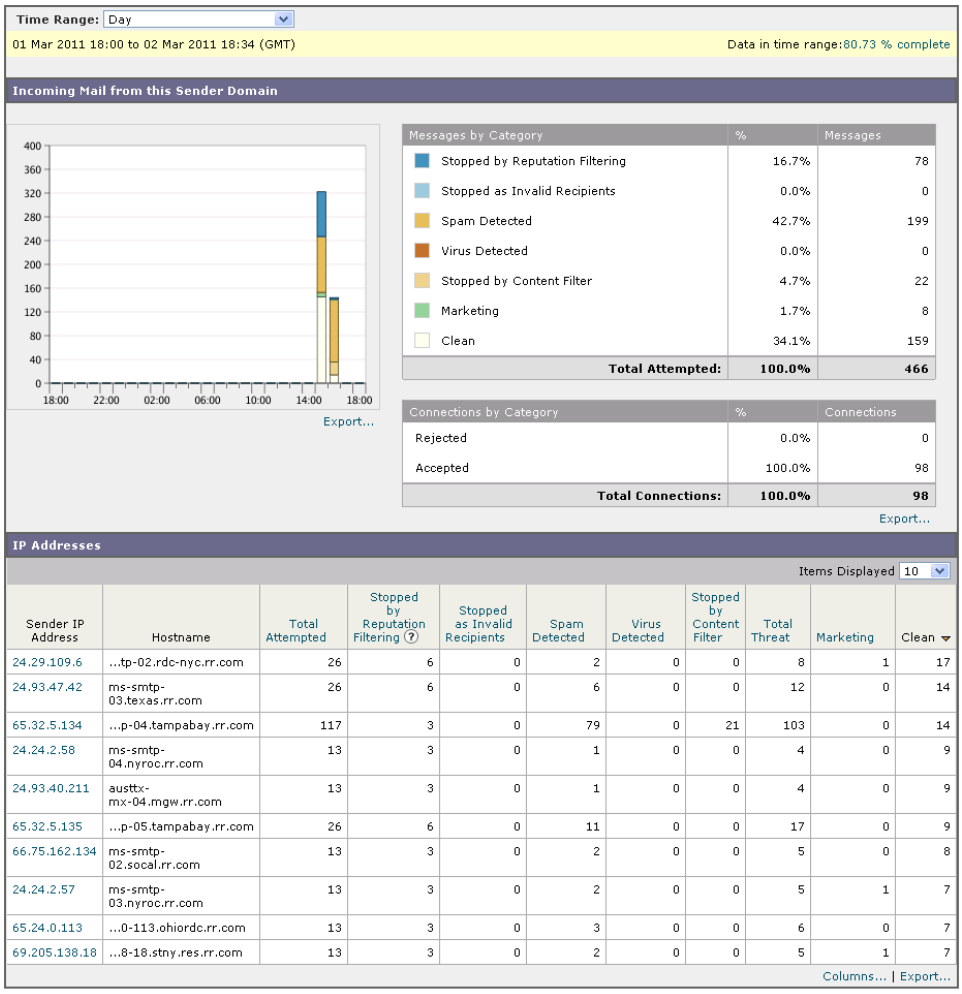

Figure 4-7 Network Owner Profile Page

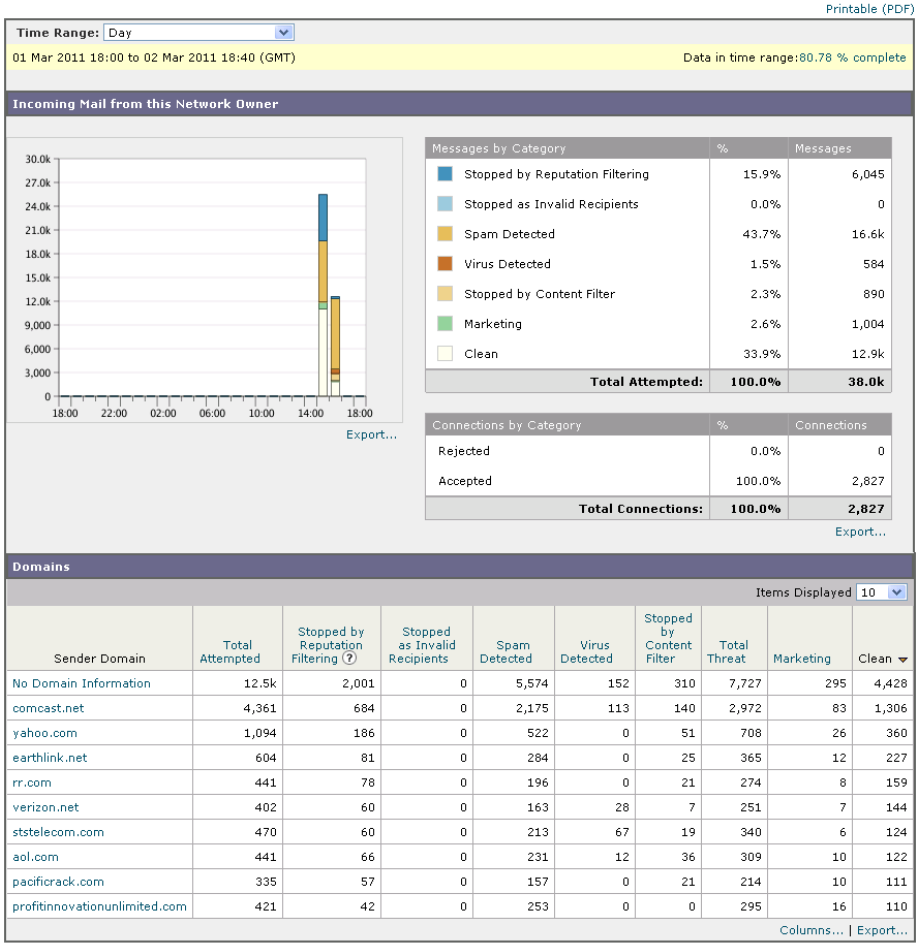

Figure 4-8 IP Address Profile Page

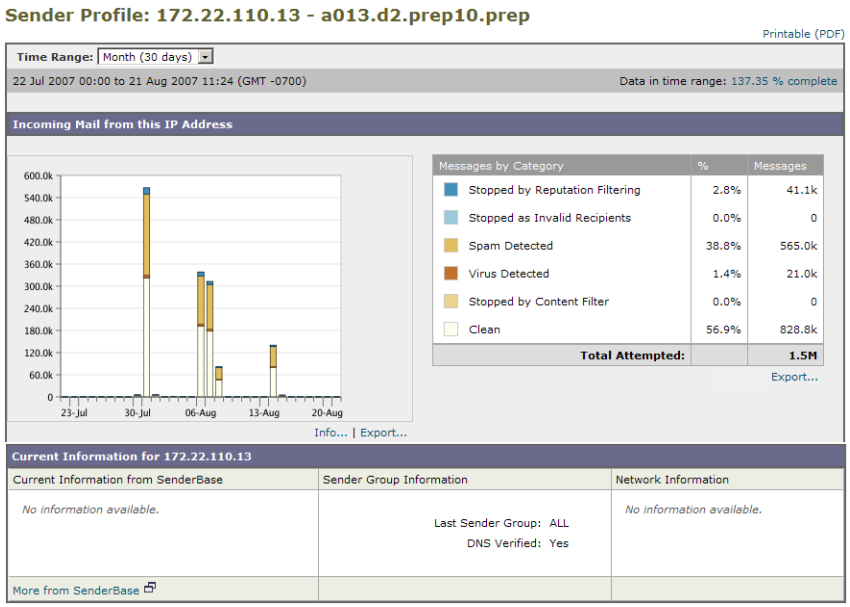

Sender Groups Report Page

The **Sender Groups repor**t 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 *Cisco IronPort AsyncOS for Email Security Configuration Guide.*

To view the Sender Groups report page, click the **Sender Groups report** link at the bottom of the Incoming Mail report page.

Figure 4-9 Sender Groups Report Page

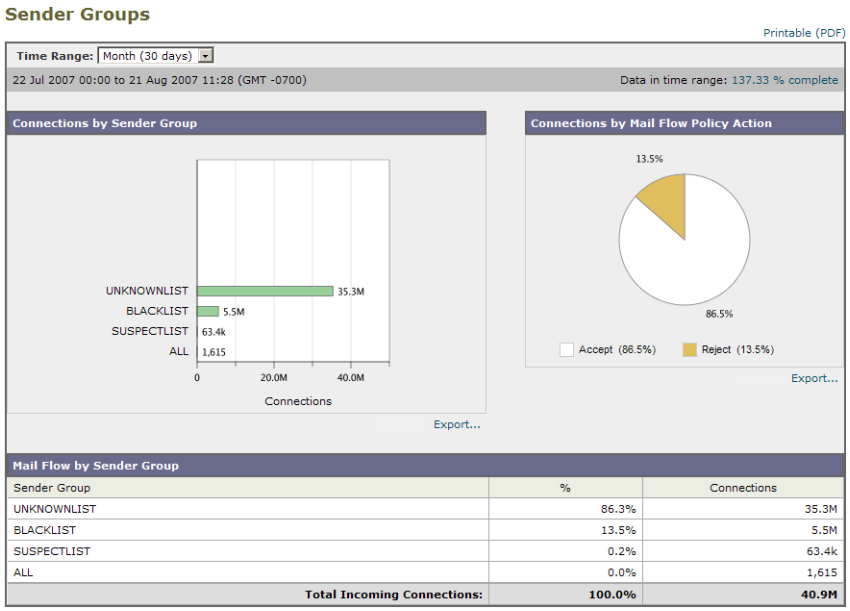

From the **Sender Group report** page you can also generate a PDF or export raw data to a CSV file. For information on printing or exporting a file, see the "Understanding the Email Reporting Pages" section on page 4-6.

Note You can generate a scheduled report for the Sender Group report page. See the "Scheduling Email Reports" section on page 4-55.

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, or stopped by a content filter?
- How many messages are delivered and how many messages are hard-bounced by the destination servers?

Figure 4-10 Email > Reporting > Outgoing Destinations Page

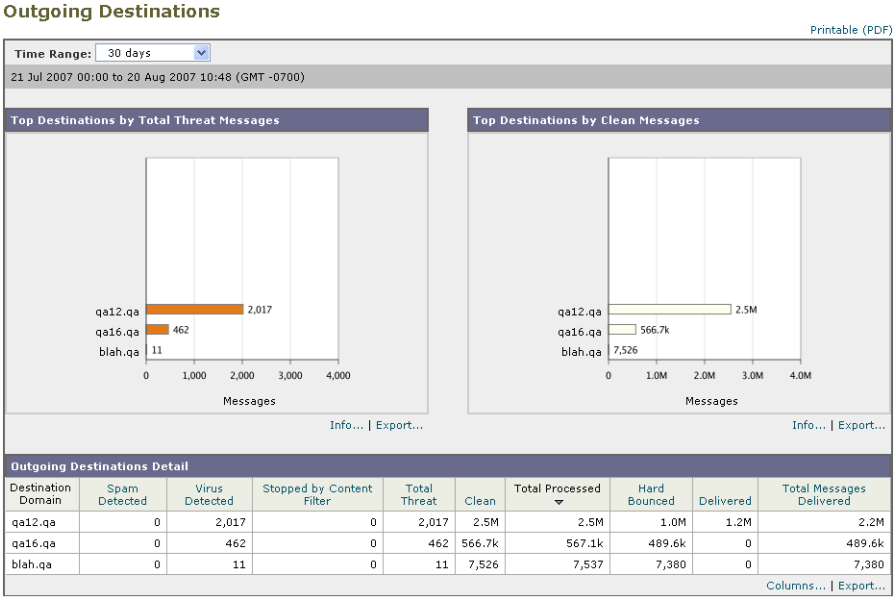

The following list explains the various sections on the **Outgoing Destinations** page:

Table 4-6 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" section on page 3-4.
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.

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" section on page 4-6.

Note You can generate a scheduled report for the Outgoing Destinations page. See the "Scheduling Email Reports" section on page 4-55.

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 or spam positive 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 Sender** page, perform the following:

Figure 4-11 Email > Reporting > Outgoing Senders Page (IP Addresses Displayed)

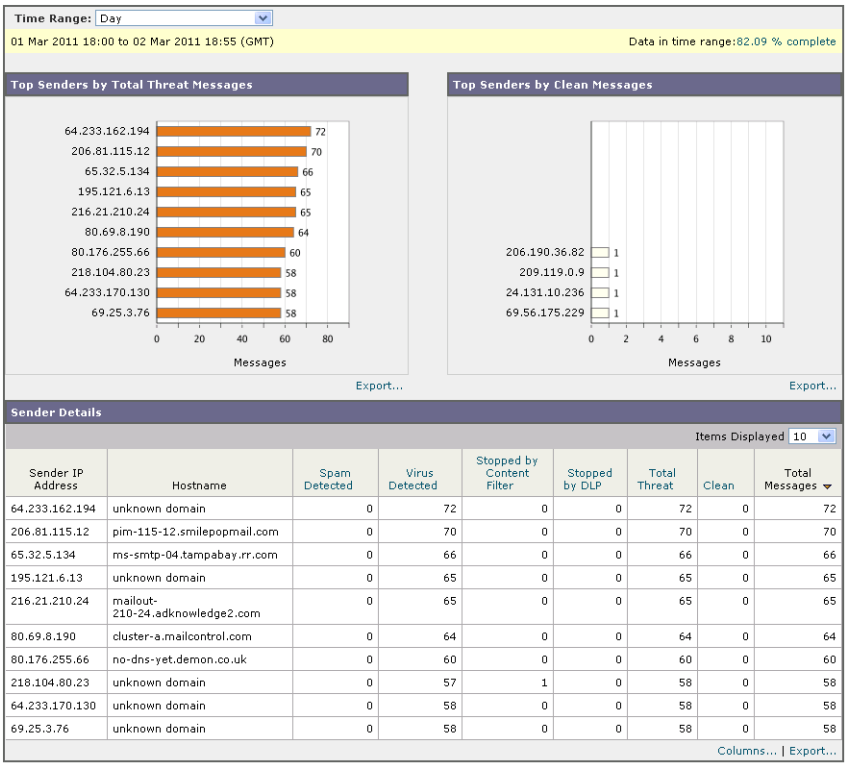

You can see the results of the Outgoing senders with two types of views:

- **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 4-7 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" section on page 3-4.
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.
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.

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" section on page 4-6.

Note You can generate a scheduled report for the **Outgoing Senders** page. See the "Scheduling Email Reports" section on page 4-55.

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 sends the most external email?
- Who receives the most clean email?
- Who receives the most spam?
- Who is triggering particular content filters?
- Are content filters stopping email from a particular user?

Figure 4-12 Email > Reporting> Internal Users Page

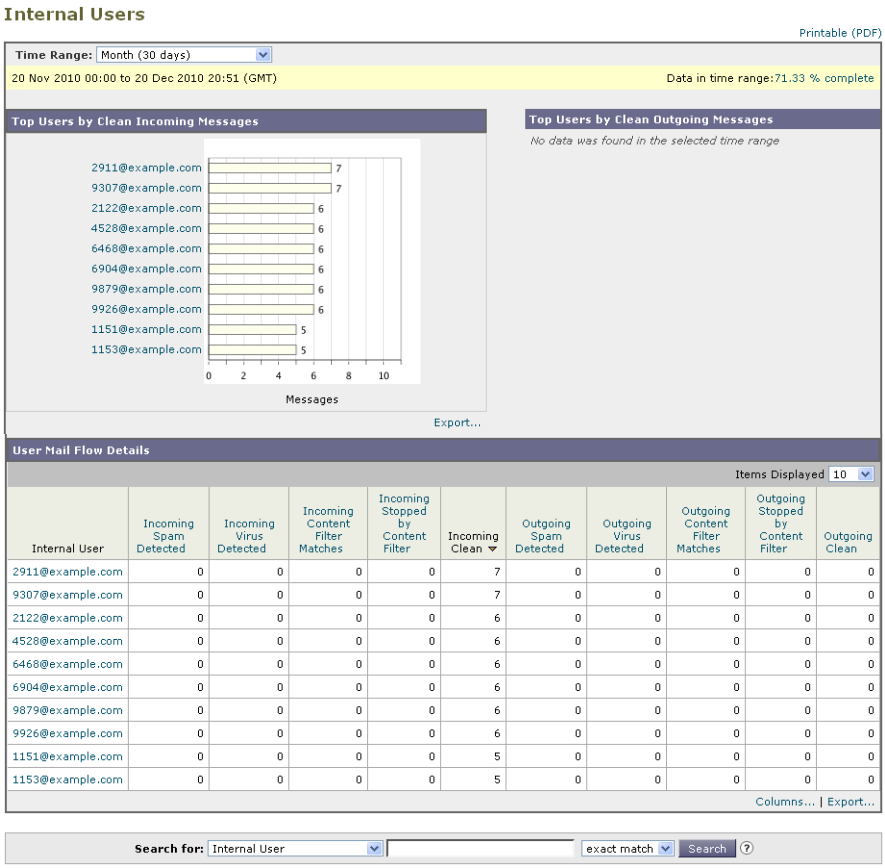

Table 4-8 Details on the Email > Reporting > Internal Users Page

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" section on page 3-4.
Top Users by Clean Incoming Messages	The top users by (by IP address or domain) of clean incoming messages sent in your organization.

Table 4-8 Details on the Email > Reporting > Internal Users Page

Section	Description
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 into Clean, Spam Detected (incoming only), Virus Detected, and Content Filter Matches. 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" section on page 4-29.

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" section on page 4-6.

Note You can generate a scheduled report for the Internal Users page. See the "Scheduling Email Reports" section on page 4-55.

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 (spam detected, virus detected, stopped by content filter, and clean). 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, page 4-32). 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 Internal Users page and the Internal User detail 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").

Figure 4-13 Internal User Search Results

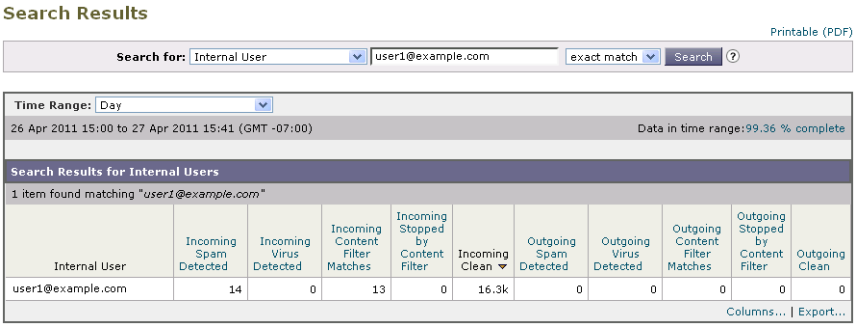

DLP Incident Summary Page

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 Cisco IronPort 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?

Figure 4-14 Email > Reporting > DLP Incidents Summary Page

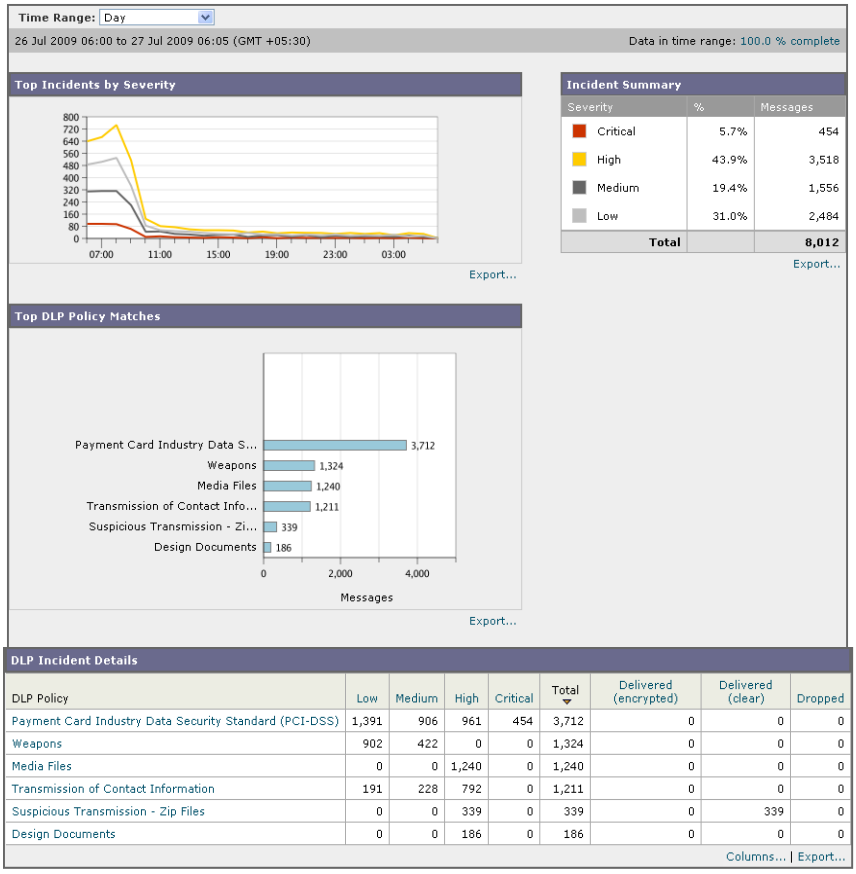

The DLP Incident Summary page contains two main sections:

- 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 4-9 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" section on page 3-4.
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.

Table 4-9 Details on the Email > Reporting > DLP Incident Summary Page

Section	Description
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" section on page 4-32.

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" section on page 4-32.

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 opens up the Internal Users page. See the "Internal Users Page" section on page 4-27 for more information.

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?

Figure 4-15 Email > Reporting > Content Filters Page

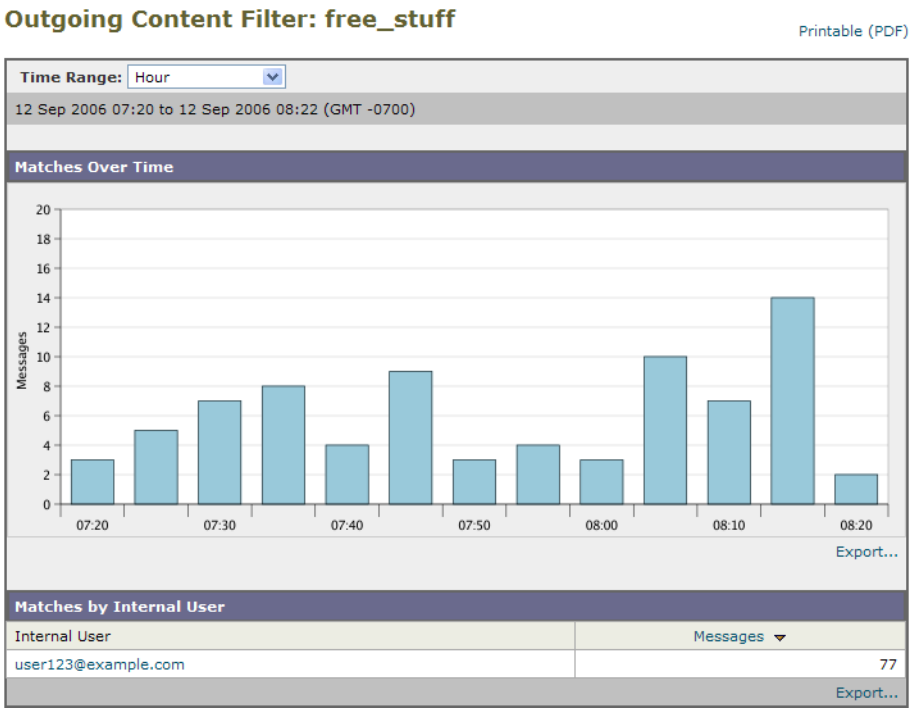

To view more information about a specific filter, click the name of the filter. The Content Filter Details page appears. For more information on Content Filter details page, see the "Content Filter Details Page" section on page 4-33.

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" section on page 4-6.

Note You can generate a scheduled report for the Content Filter page. See the "Scheduling Email Reports" section on page 4-55.

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, page 4-29.

Figure 4-16 Content Filters Details Page

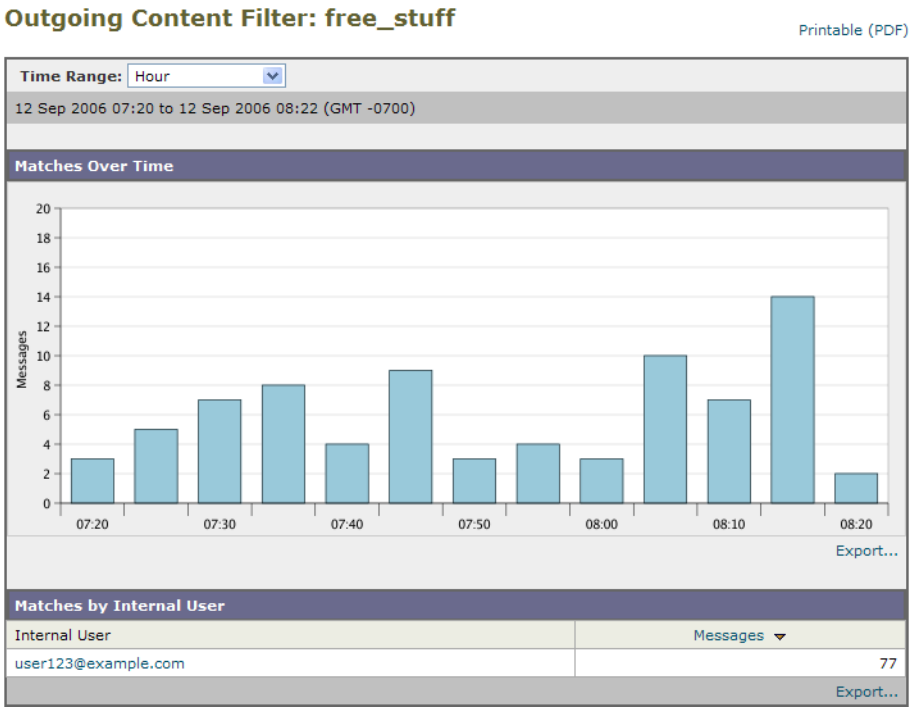

Virus Types Page

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.

Figure 4-17 Email > Reporting > Virus Types Page

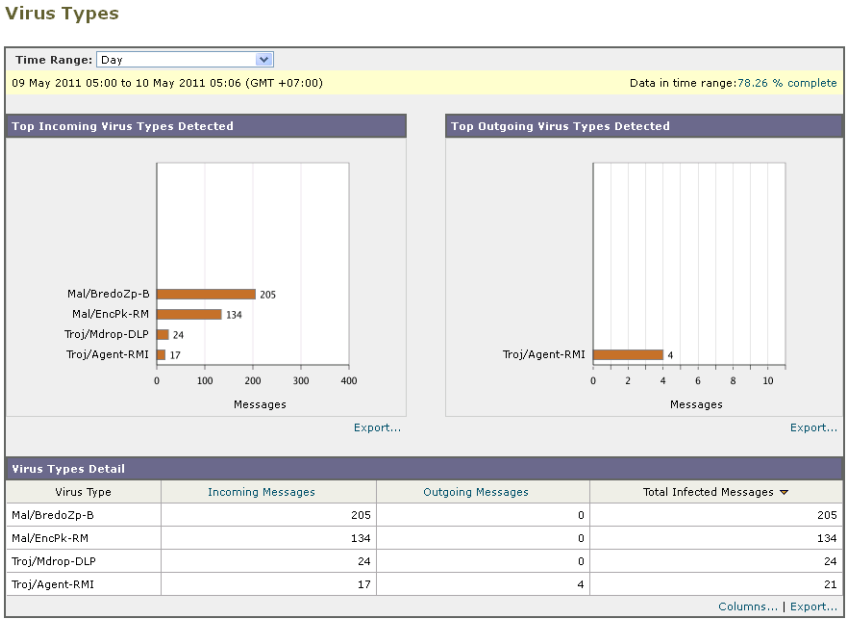

If you run multiple virus scanning engines, the Virus Types page includes results from all enabled virus scanning engines. The name of the virus that appears on the page is determined by the virus scanning engines. If more than one scanning engine detects a virus, it is possible to have more than one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

Table 4-10 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" section on page 3-4.
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" section on page 4-6.

Note You can generate a scheduled report for the **Virus Types** page. See the "Scheduling Email Reports" section on page 4-55.

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?

The TLS Connections Report page has the following sections:

- TLS Connections Report Page: Incoming Connections
- TLS Connections Report Page: Outgoing Connections

Figure 4-18 TLS Connections Report Page: Incoming Connections

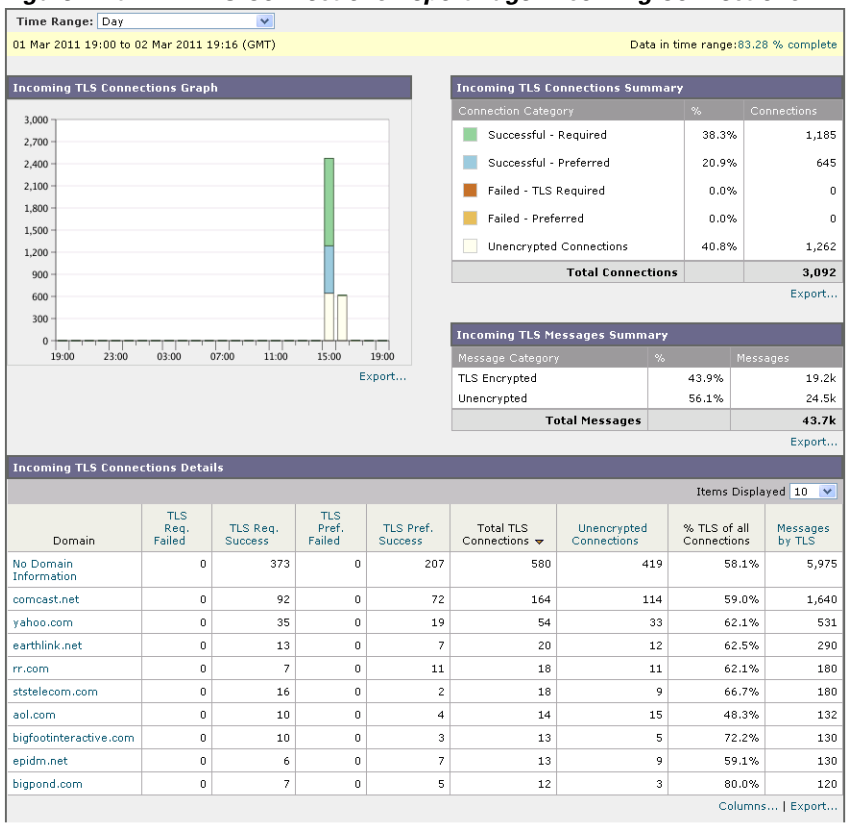

Figure 4-19 TLS Connections Report Page: Outgoing Connections

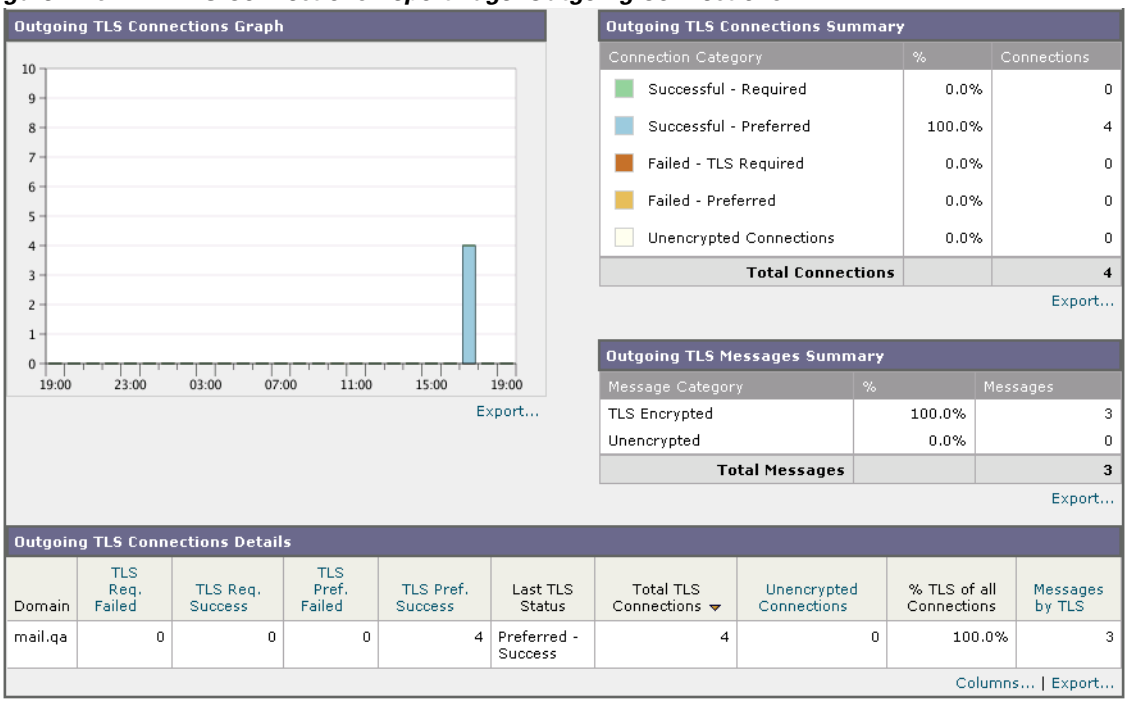

Table 4-11 Details on the Email > Reporting > TLS Connections Page

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" section on page 3-4.
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.

Table 4-11 Details on the Email > Reporting > TLS Connections Page

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 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.

Figure 4-20 Inbound SMTP Authentication Page

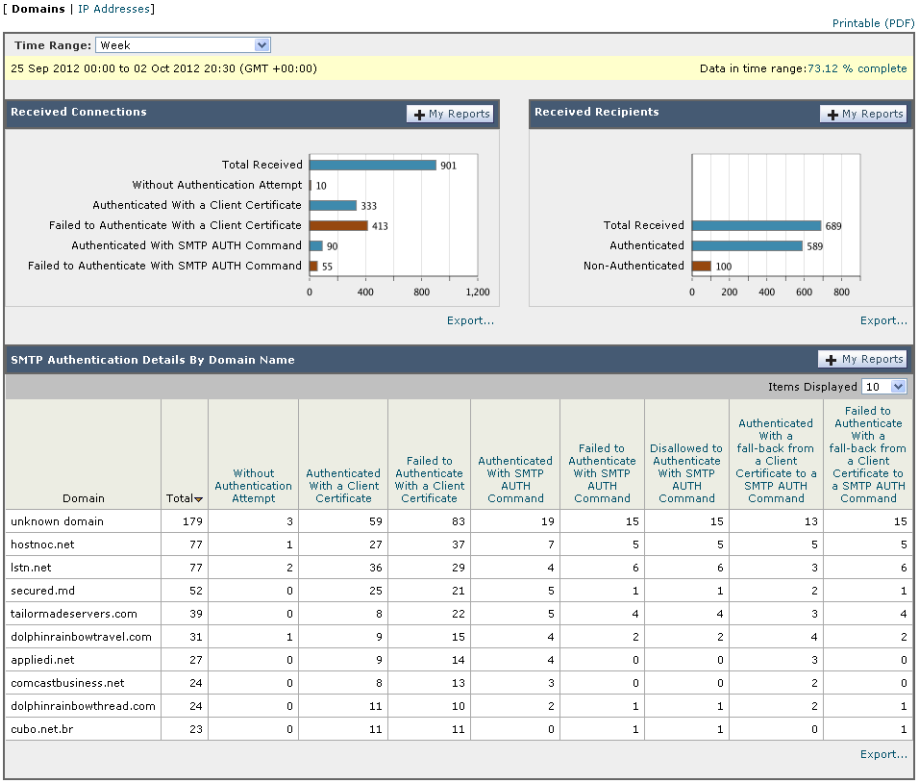

Rate Limits Page

Rate Limiting by envelope sender allows you to limit the number of email messages per time interval from an individual sender, based on the mail-from address. The **Email > Reporting > Rate Limits** page allows you to quickly identify individual senders of large numbers of messages. Use this report to help you to:

- Control spam from internal user accounts, for example in cases when a user's credentials are compromised and the account is used to send spam in bulk.
- Identify compromised user accounts.
- Limit out-of-control applications that use email for notifications, alerts, automated statements, etc.
- Avoid damaging your organization’s online reputation and the attendant hassles resulting from this situation

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 *Cisco IronPort AsyncOS for Email Security Configuration Guide*.

Figure 4-21 Rate Limits Page

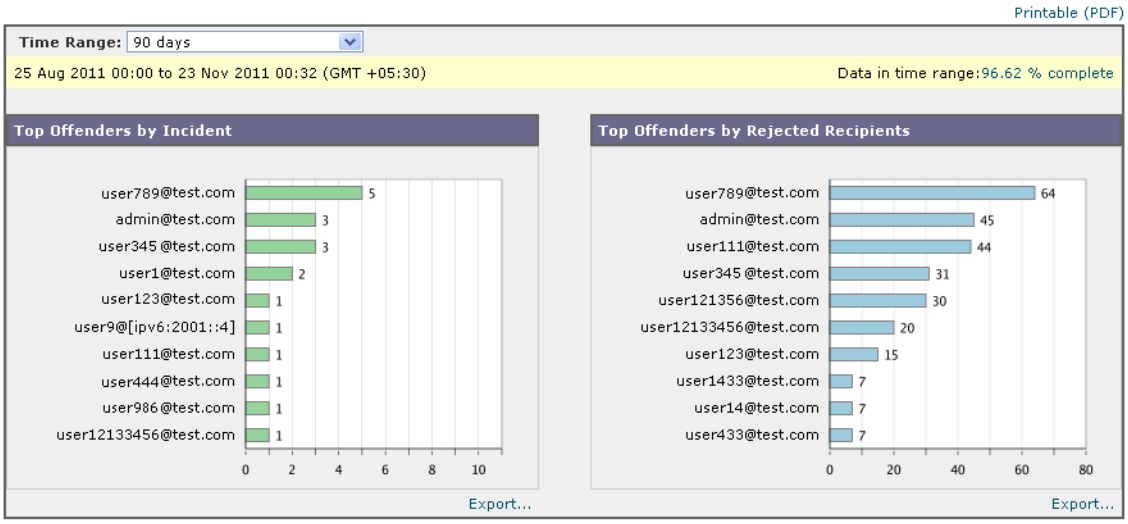

Outbreak Filters Page

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?

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 Cisco IronPort appliance. Local outbreak data does not include non-viral threats. Global outbreak data represents all outbreaks detected by the Cisco IronPort 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 Cisco IronPort Threat Operations Center and the release of an anti-virus signature by a major vendor. Note that not every global outbreak affects your Cisco IronPort appliance. A value of "--" indicates either a protection time does not exist, or the signature times were not available from the anti-virus vendors (some vendors may not report signature times). This does not indicate a protection time of zero, rather it means that the information required to calculate the protection time is not available.

The Quarantined Messages section summarizes Outbreak Filters quarantining, and is a useful gauge of how many potential threat messages Outbreak Filters are catching. Quarantined messages are counted at time of release. Typically, messages will be quarantined before anti-virus and anti-spam rules are available. When released, they will be scanned by the anti-virus and anti-spam software and determined

to be positive or clean. Because of the dynamic nature of Outbreak tracking, the rule under which a message is quarantined (and even the associated outbreak) may change while the message is in the quarantine. Counting the messages at the time of release (rather than the time of entry into the quarantine) avoids the confusion of having counts that increase and decrease.

The Threat Details listing displays information about specific outbreaks, including the threat category (virus, scam, or phishing), threat name, a description of the threat, and the number of messages identified. For virus outbreaks, the Past Year Virus Outbreaks include the Outbreak name and ID, time and date a virus outbreak was first seen globally, the protection time provided by Outbreak filters, and the number of quarantined messages. You can select either global or local outbreaks as well as the number of messages to display via the menu on the left. You can sort the listing by clicking on the column headers.

The First Seen Globally time is determined by the Cisco IronPort Threat Operations Center, based on data from the Cisco IronPort 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 Cisco IronPort 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.

Figure 4-22 Outbreaks Page

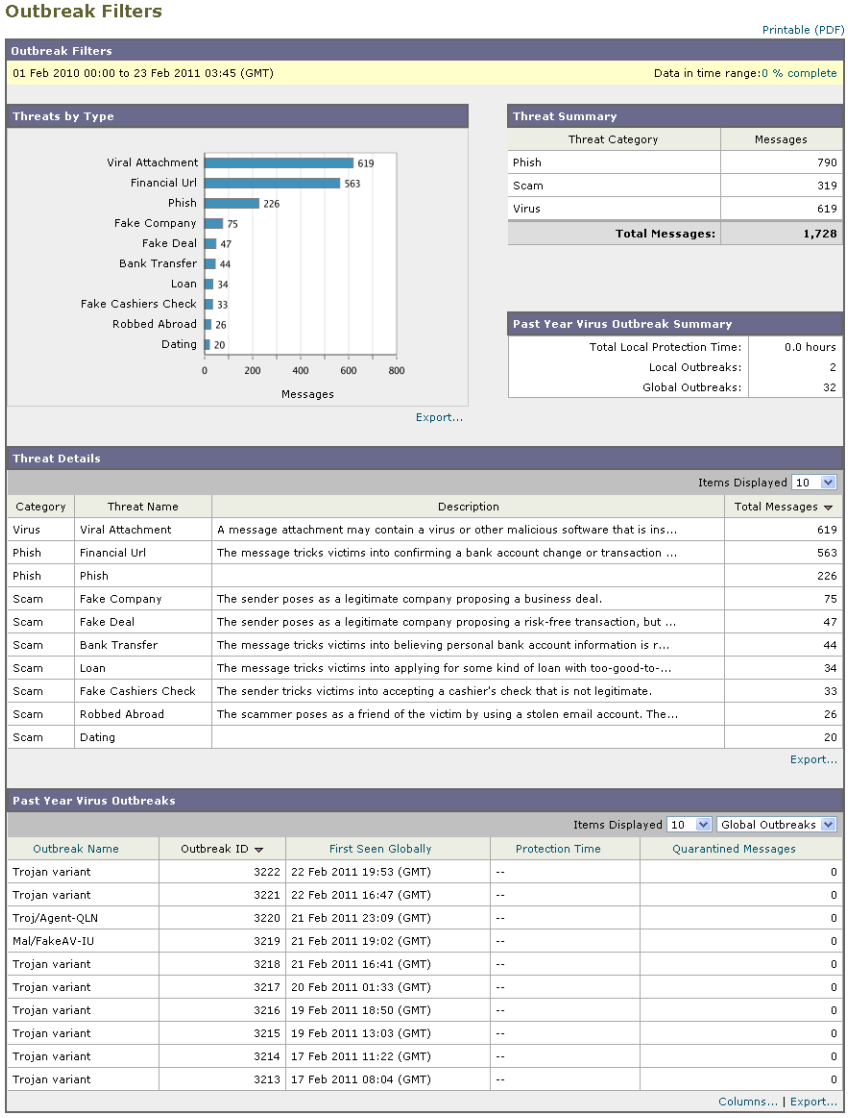

Note The Security Management appliance must be able to communicate with `downloads.cisco.com` to correctly populate the tables on the Outbreak Filters page.

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, page 4-45 and System Capacity – Outgoing Mail, page 4-46.
- **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, page 4-44.
- **Resource Conservation Mode:** When a Cisco IronPort 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 Cisco IronPort 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. RCM is not tracked by the System Capacity page.

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 System Capacity – Workqueue page shows the volume of messages in work queues over a specified time period. It also shows the maximum messages in work queues over the same time period. You can view data for a day, week, month, or year. Occasional spikes in the Workqueue graphs are normal and expected. If the spikes occur with increasing frequency and are maintained over a long period of time, this may indicate a capacity issue. When reviewing the work queue page, you may want to measure the frequency of work queue backups, and take note of work queue backups that exceed 10,000 messages.

Figure 4-23 System Capacity – Workqueue

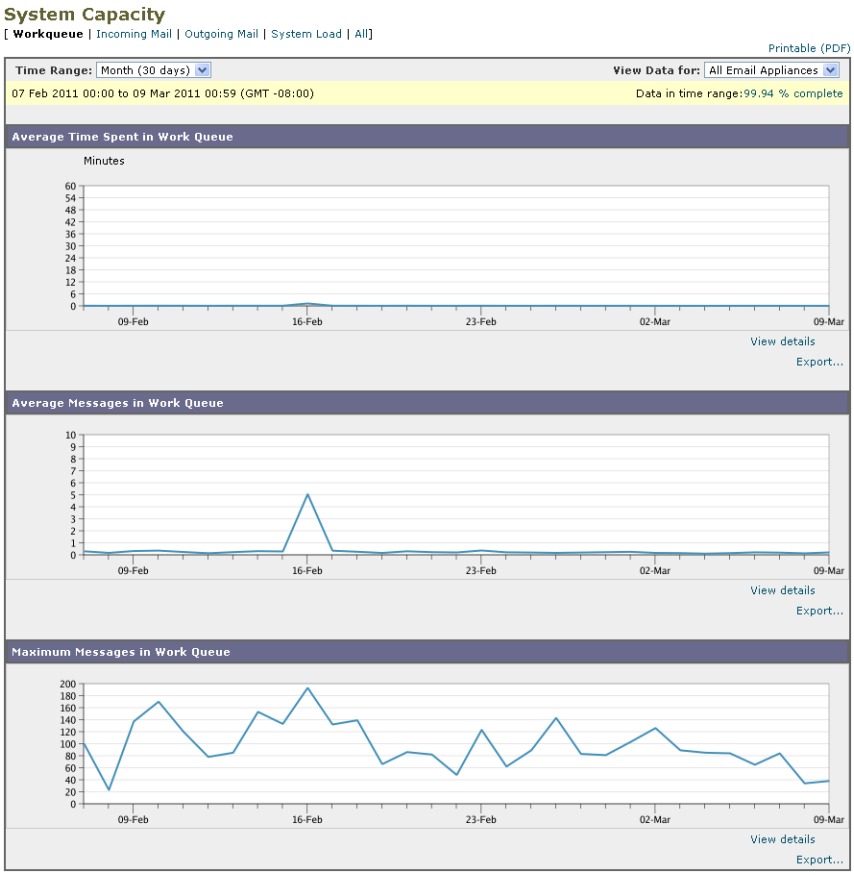

System Capacity – Incoming Mail

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.

Figure 4-24 System Capacity – Incoming Mail

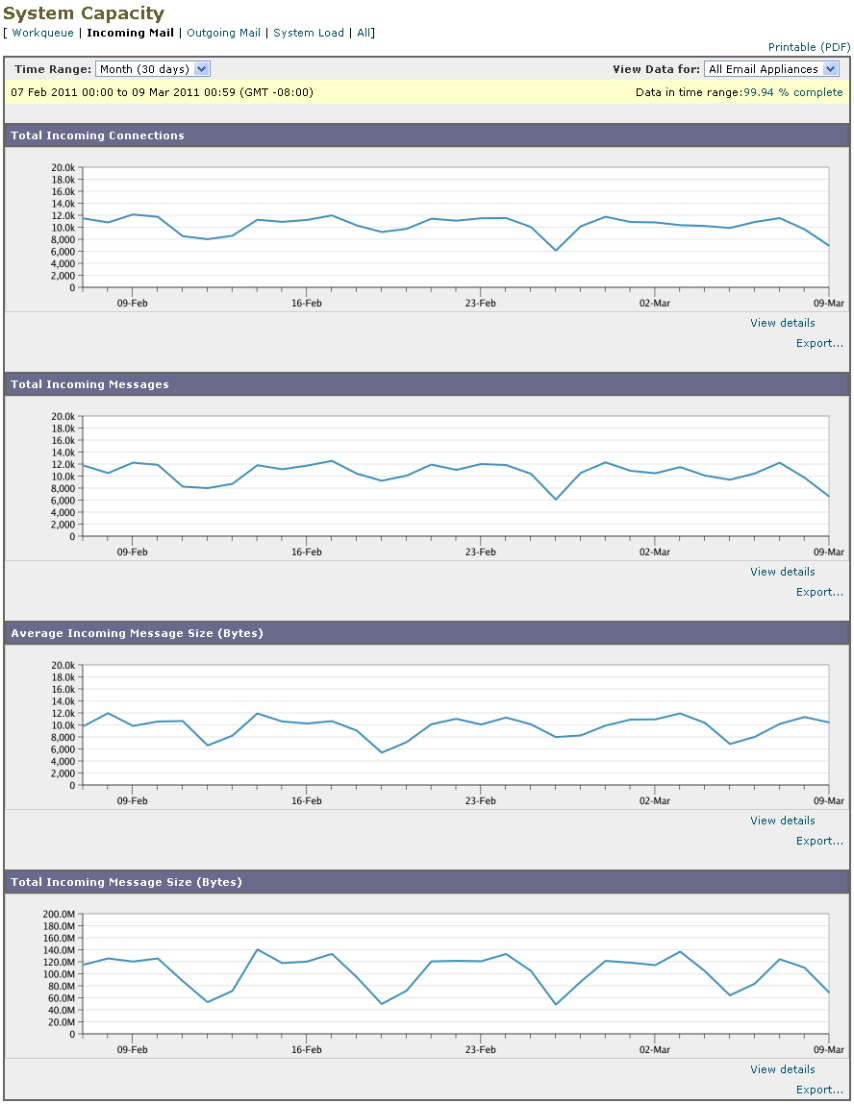

System Capacity – Outgoing Mail

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.

Figure 4-25 System Capacity – Outgoing Mail

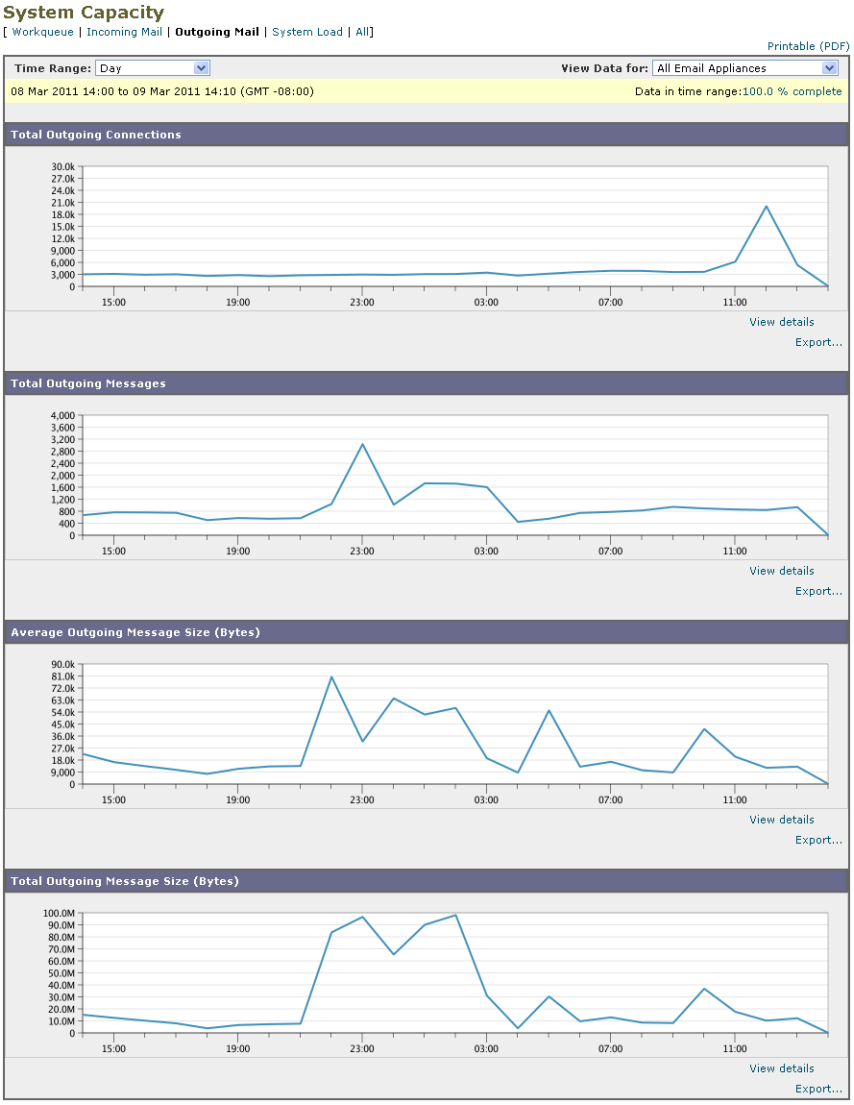

System Capacity – System Load

The system load report shows the overall CPU usage on the Email Security appliances. AsyncOS is optimized to use idle CPU resources to improve message throughput. High CPU usage may not indicate a system capacity problem. If the high CPU usage is coupled with consistent, high-volume memory page swapping, you may have a capacity problem. This page also shows a graph that displays the amount of CPU used by different functions, including mail processing, spam and virus engines, reporting, and quarantines. The CPU-by-function graph is an indicator of which areas of the product use the most resources on your system. If you need to optimize your appliance, this graph can help you determine which functions may need to be tuned or disabled.

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

Figure 4-26 System Capacity – System Load

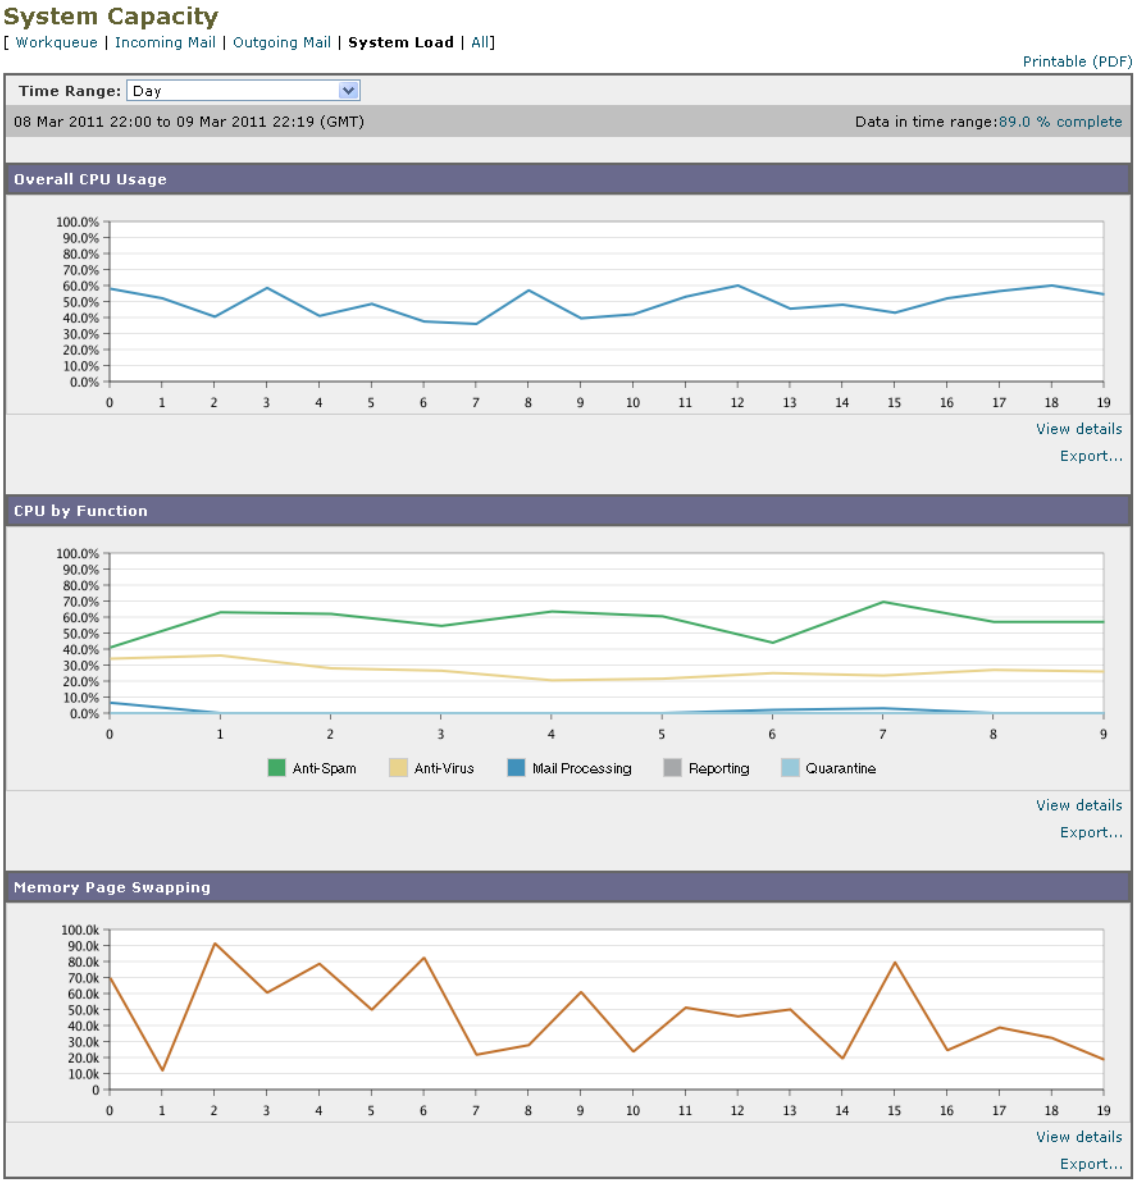

Note About Memory Page Swapping

The system is designed to swap memory regularly, so some memory swapping is expected and is not an indication of problems with your appliance. Unless the system *consistently* swaps memory in high volumes, memory swapping is normal and expected behavior (especially on C150 appliances). For example, Figure 4-27 shows a system that consistently swaps memory in high volumes. To improve performance, you may need to add Cisco IronPort appliances to your network or tune your configuration to ensure maximum throughput.

Figure 4-27 System Capacity – System Load (System Under Heavy Load)

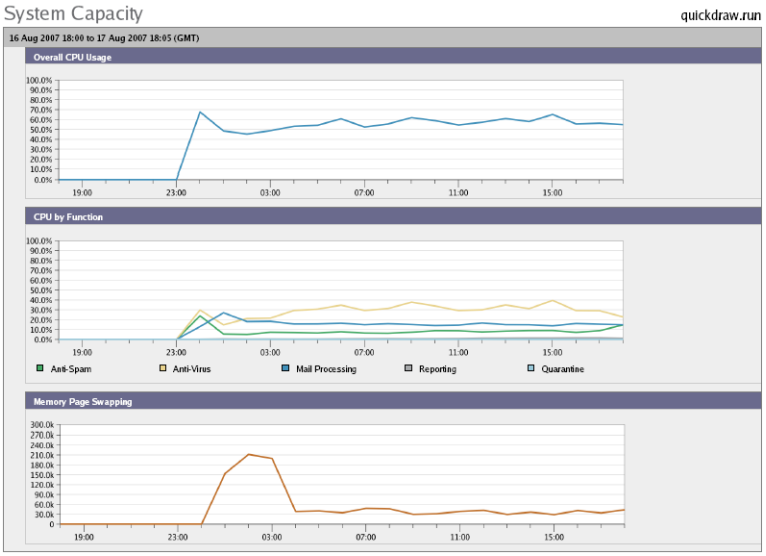

System Capacity – All

The **All** page consolidates all the previous system capacity reports onto a single page so you can view the relationship between the different reports. For example, you might 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).

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.

Figure 4-28 Email Reporting Data Availability Page

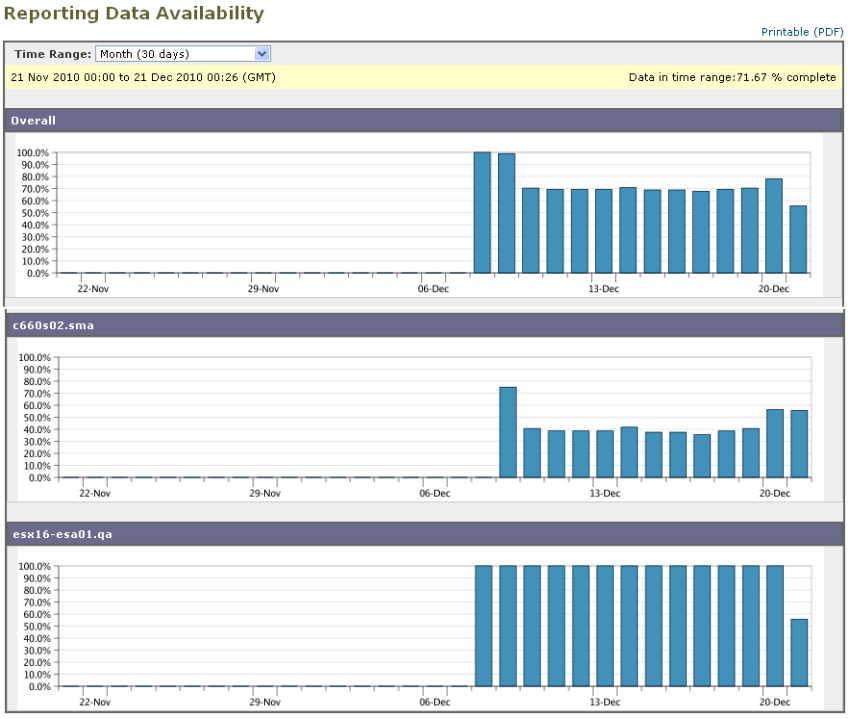

All data resource utilization and email traffic trouble spots are shown from this page, including the data availability for the overall appliances that are managed by the Security Management appliance.

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" section on page 4-32.
- DLP Incident Summary—For information on what is included on this page, see the "DLP Incident Summary Page" section on page 4-30.
- 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 *Cisco IronPort AsyncOS for Email Security Daily Management Guide.*
- Domain-Based Executive Summary—This report is based on the Email Reporting Overview Page, and is limited to a group of specified domains. For information on what is included, see the "Domain-Based Executive Summary Report" section on page 4-52.

- Executive Summary—This report is based on the information from the Email Reporting Overview Page. For information on what is included, see the "Domain-Based Executive Summary Report" section on page 4-52.
- Incoming Mail Summary — For information on what in included on this page, see the "Incoming Mail Page" section on page 4-14.
- Internal Users Summary—For information on what is included on this page, see the "Internal Users Page" section on page 4-27.
- Outbreak Filters—For information on what is included on this page, see the "Outbreak Filters Page" section on page 4-41.
- Outgoing Destinations—For information on what is included on this page, see the "Outgoing Destinations Page" section on page 4-24.
- Outgoing Mail Summary—For information on what is included on this page, see the "Outgoing Senders Page" section on page 4-26.
- Outgoing Senders: Domains —For information on what is included on this page, see the "Outgoing Senders Page" section on page 4-26.
- Sender Groups —For information on what is included on this page, see the "Sender Groups Report Page" section on page 4-23.
- System Capacity—For information on what is included on this page, see the "System Capacity Page" section on page 4-43.
- TLS Connections—For information on what is included on this page, see the "TLS Connections Page" section on page 4-36.
- Virus Types—For information on what is included on this page, see the "Virus Types Page" section on page 4-34.

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 and Tracking Data, page 3-8.

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, page 4-58.

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
- Executive Summary Reports

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. If multiple domains are specified, the appliance aggregates the data for all those domains into a single report. Unlike other scheduled reports, Domain-Based Executive Summary reports are not archived.

Because messages blocked by 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 Cisco IronPort 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 Cisco IronPort Email Security appliance. However, this option can impact system performance. For more information about delayed HAT rejections, see the documentation for Cisco IronPort AsyncOS for Email Security.

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.

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.

To create a Domain-Based Executive Summary report:

Procedure

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

To schedule the report:

a. Choose **Email > Reporting > Scheduled Reports**.

b. Click **Add Scheduled Report**.

To create an on-demand report:

a. Choose **Email > Reporting > Archived Reports.**

b. Click **Generate Report Now**.

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

Figure 4-29 Adding a Domain-Based Executive Summary Report

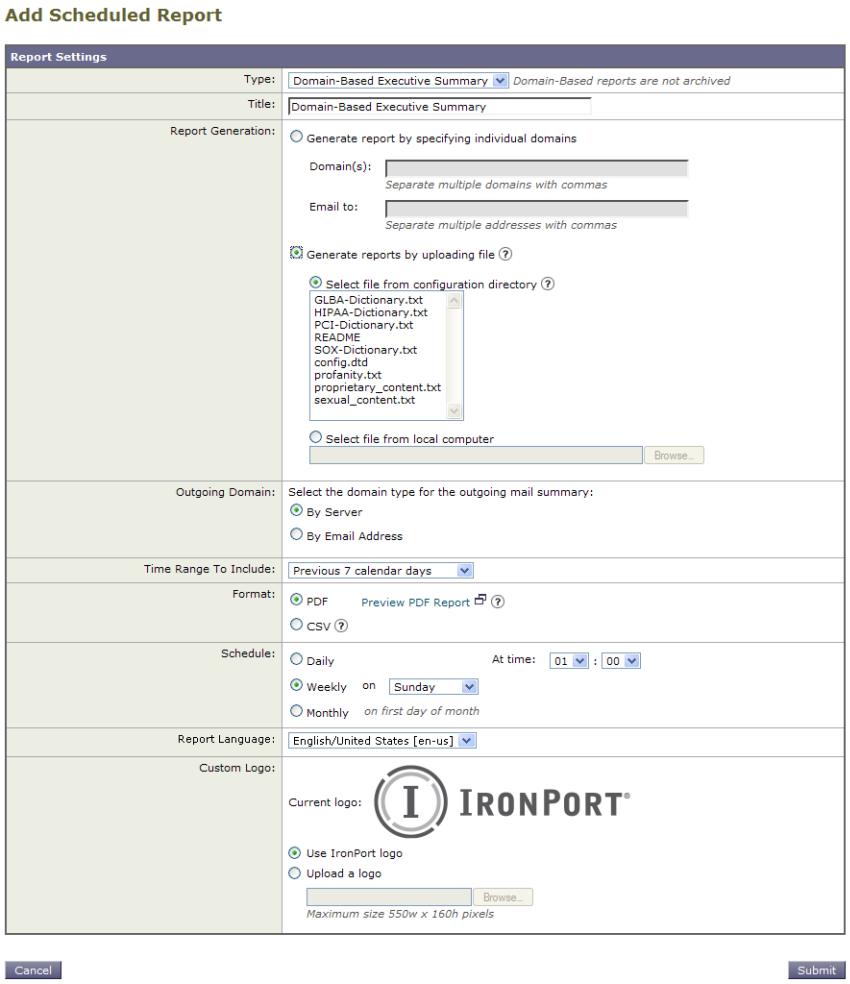

Step 3 Specify the domains to include in the report and the email addresses for the report recipients. You can select one of the following options for generating the report:

- **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 Domain-Based Executive Summary Report Configuration Files, page 4-54.

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 IronPort 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 and Tracking Data, page 3-8.

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

Domain-Based Executive Summary Report Configuration Files

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.

Executive Summary Reports

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 Email Reporting Overview Page. For more information on the Email Reporting Overview page, see "Email Reporting Overview Page" section on page 4-10.

Scheduling Email Reports

You can schedule any of the reports listed in About Scheduled and On-Demand Email Reports, page 4-50.

To manage report scheduling, see the following:

- Adding Scheduled Reports, page 4-55
- Editing Scheduled Reports, page 4-56
- Discontinuing Scheduled Reports, page 4-57

Adding Scheduled Reports

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

Procedure

Step 1 On the Security Management appliance window, choose **Email > Reporting > Scheduled Reports**.

Step 2 Click **Add Scheduled Report**.

Step 3 Choose your report type.

For descriptions of the report types, see About Scheduled and On-Demand Email Reports, page 4-50.

Note For information about the settings for a Domain-Based Executive Summary report, see Domain-Based Executive Summary Report, page 4-52.

Note 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 4 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 5 Choose the time range for the report from the **Time Range to Include** drop-down menu.

Step 6 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 7 Depending on the report, for **Number of Rows**, choose the amount of data to include.

Step 8 Depending on the report, choose the column by which to sort the report.

Step 9 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 10 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 11 Choose a language for the report.

For Asian languages, see important information at Printing and Exporting Reporting and Tracking Data, page 3-8.

Step 12 Click **Submit**.

Editing Scheduled Reports

Procedure

Step 1 On the Security Management appliance window, choose **Email > Reporting > Scheduled Reports**.

Step 2 Click the report name link in the Report Title column that you want to modify.

Step 3 Modify the report settings.

Step 4 Submit and commit your changes.

Discontinuing Scheduled Reports

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

Procedure

Step 1 On the Security Management appliance window, choose **Email > Reporting > Scheduled Reports**.

Step 2 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 3 Click **Delete**.

Note Any archived versions of deleted reports are *not* automatically deleted. To delete previously-generated reports, see Deleting Archived Reports, page 4-59.

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, page 4-6, you can save PDFs or raw-data CSV files for the reports listed in About Scheduled and On-Demand Email Reports, page 4-50 at any time, for the time frame that you specify.

To generate an on-demand report perform the following:

Procedure

Step 1 On the Security Management appliance, choose **Email > Reporting > Archived Reports**.

Step 2 Click **Generate Report Now**.

Step 3 Choose a report type.

For descriptions of the report types, see About Scheduled and On-Demand Email Reports, page 4-50.

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.

Note For information about the settings for a Domain-Based Executive Summary report, see Domain-Based Executive Summary Report, page 4-52.

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

Note the custom time range option.

Step 6 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 7 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 8 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 9 Select a language for this report. For generating PDFs in Asian languages, see important information at Printing and Exporting Reporting and Tracking Data, page 3-8.

Step 10 Click **Deliver This Report** to generate the report.

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 12 instances of each scheduled report, up to 1000 total versions for all reports. The limit of 12 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 Appendix A, "IP Interfaces and Accessing the Appliance" 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.

Procedure

Step 1 Choose **Email > Reporting > Archived Reports**.

The list of Archived Reports appears.

Figure 4-30 Archived Reports

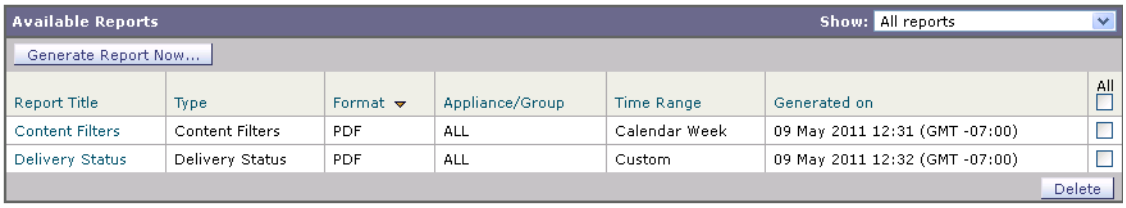

Step 2 To locate a particular report if the list is long, filter the list by choosing a report type from the **Show** menu, or click a column heading to sort by that column.

Step 3 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, page 4-58. However, you can manually delete unneeded reports.

To manually delete Archived reports, perform the following:

Procedure

Step 1 On the Security Management appliance, choose **Email > Reporting> Archived Reports**.

The Archived reports that are available are displayed.

Step 2 Select the checkbox for one or more reports to delete.

Step 3 Click **Delete**.

Step 4 To prevent future instances of scheduled reports from being generated, see Discontinuing Scheduled Reports, page 4-57.

CHAPTER 5

Using Centralized Web Reporting and Tracking

- Centralized Web Reporting Overview, page 5-1
- Setting Up Centralized Web Reporting, page 5-2
- Working with Interactive Web Reporting Pages, page 5-7
- Understanding the Web Reporting Pages, page 5-7
- About Scheduled and On-Demand Web Reports, page 5-61
- Scheduling Web Reports, page 5-61
- Generating Web Reports on Demand, page 5-65
- Viewing and Managing Archived Web Reports, page 5-66

Centralized Web Reporting Overview

The Web reporting functionality aggregates information from individual security features and records data that can be used to monitor your web 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.

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

Domain Information

For a domain, the web reporting feature can generate the following data elements to be on a domain report. For example, if you are generating a report on the Facebook.com domain, the report may contain:

- A list of the top users who accessed Facebook.com
- A list of the top URLs that were accessed within Facebook.com

User

For a user, the web reporting feature can generate data elements to be on a user report. For example, for the user report titled ‘Jamie’, the report may contain:

- A list of the top domains that the user ‘Jamie’ accessed
- A list of the top URLs that were malware or virus positive

- A list of the top categories that the user ‘Jamie’ accessed

Category

For a category, the web reporting feature can generate data to be included in a category report. For example, for the category ‘Sports’, the report may contain:

- A list of the top domains that were in the ‘Sports’ category
- A list of the top users who accessed the ‘Sports’ category

In all of these examples, these reports are intended to give a comprehensive view about a particular item on the network so that the administrator can take action.

General

For a detailed description on logging pages versus reporting pages, see the “Logging Versus Reporting” section on page 14-1.

Note You can retrieve all the domain information that a user goes to, not necessarily the specific URL that is accessed. For information on a specific URL that the user is accessing, what time they went to that URL, whether that URL is allowed, etc., use the Searching for Transactions Processed by Web Proxy Services on the Web Tracking page.

Note The Web Security appliance only stores data if local reporting is used. If centralized reporting is enabled for the Web Security appliance then the Web Security appliance retains ONLY System Capacity and System Status data. If Centralized Web Reporting is not enabled, the only reports that are generated are System Status and System Capacity.

There are several ways that you can view web reporting data on the Security Management appliance.

- To view interactive report pages, see Understanding the Web Reporting Pages, page 5-7.
- To generate a report on demand, see Generating Web Reports on Demand, page 5-65.
- To schedule generation of reports on a regular, recurring basis, see About Scheduled and On-Demand Web Reports, page 5-61.
- To view archived versions of previously run reports (both scheduled and generated on demand), see Viewing and Managing Archived Web Reports, page 5-66.

Setting Up Centralized Web Reporting

To set up centralized web reporting, complete the following steps in order:

- Enabling Centralized Web Reporting on the Security Management Appliance, page 5-3
- Enabling Centralized Web Reporting on Web Security Appliances, page 5-3
- Adding the Centralized Web Reporting Service to Each Managed Web Security Appliance, page 5-3
- Anonymizing User Names in Web Reports, page 5-4

Enabling Centralized Web Reporting on the Security Management Appliance

Before you use web reporting on the Security Management appliance, you must enable web reporting on the Security Management appliance.

Procedure

Step 1 Before enabling centralized web reporting, ensure that sufficient disk space is allocated to that service. See Managing Disk Usage, page 13-50.

Step 2 On the Security Management appliance, choose **Management Appliance > Centralized Services > Web > Centralized Reporting**.

Step 3 If you are enabling centralized reporting 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 enabling centralized reporting after it has previously been disabled:

a. Click **Edit Settings**.

b. Select the **Enable Centralized Web Report Services** checkbox.

c. You can address Anonymizing User Names in Web Reports, page 5-4 now or later.

Step 5 Submit and commit your changes.

Note If you have enabled web reporting on the appliance, and there is no disk space allocated for this action, centralized web reporting will not work until disk space is allocated. As long as the quota you are setting the Web Reporting and Tracking to is larger than the currently used disk space, you will not lose any Web Reporting and Tracking data. See the "Managing Disk Usage" section on page 13-50, for more information.

Enabling Centralized Web Reporting on Web Security Appliances

All Web Security appliances should be configured and working as expected before you enable centralized reporting.

You must enable centralized reporting on each Web Security appliance that will use centralized reporting.

See the "Enabling Centralized Reporting" section in the *Cisco IronPort AsyncOS for Web Security User Guide*.

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

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

Procedure

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

Step 2 If you have already added the Web Security appliance to the list:

a. Click the name of a Web Security Appliance.

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

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

a. Click **Add Web 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 Cisco IronPort 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 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 Web Security Appliance for which you want to enable Centralized Reporting.

Step 6 Commit your changes.

Anonymizing User Names in Web Reports

By default, user names appear on reporting pages and PDFs. However, to protect user privacy, you may want to make user names unrecognizable in web reports.

Note Users with Administrator privileges on this appliance can always see user names when viewing interactive reports.

Figure 5-1 Reporting Page Showing User Names

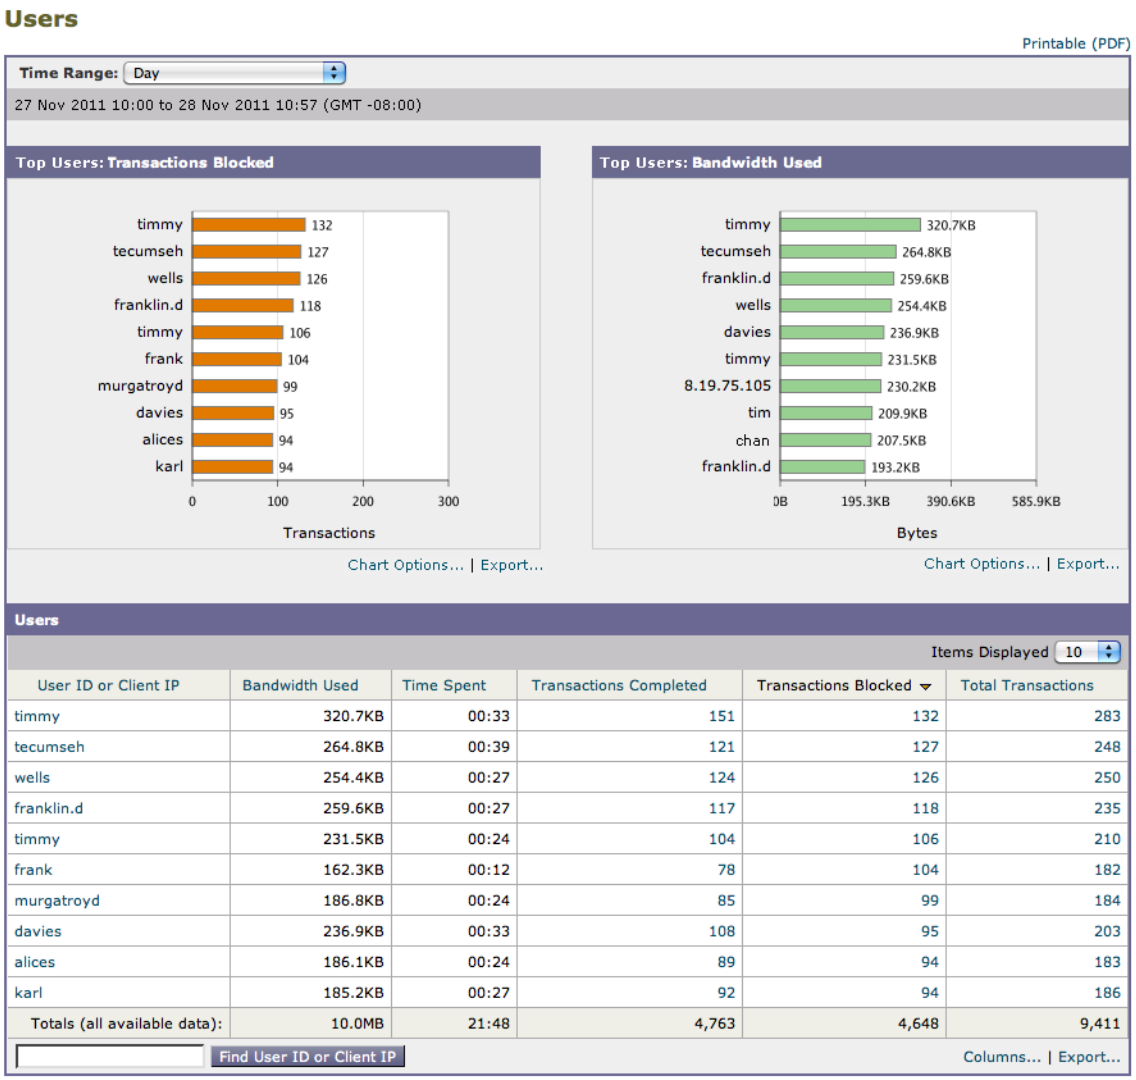

Figure 5-2 Reporting Page With Anonymized User Names

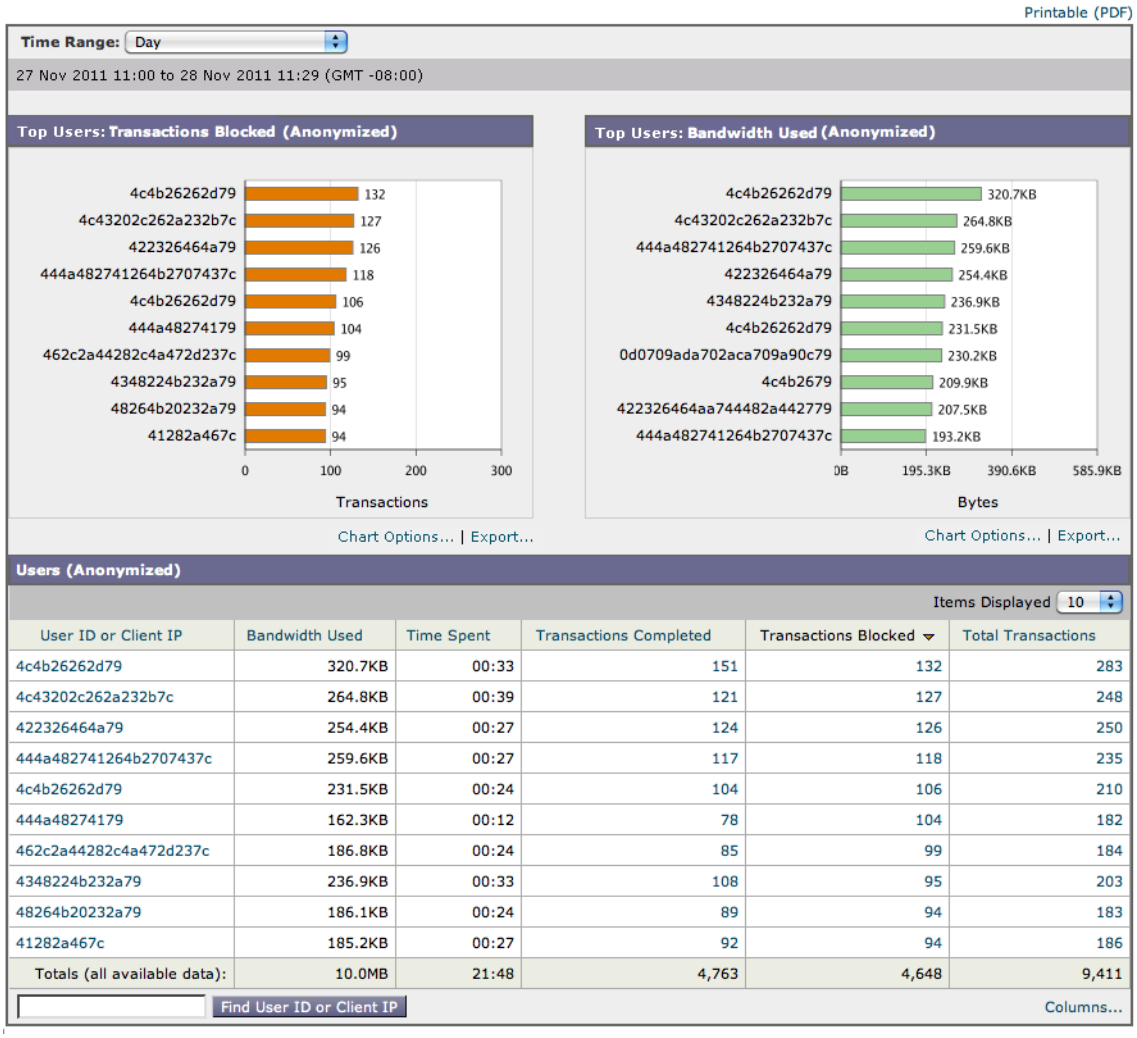

To make user names unrecognizable in reports:

Procedure

Step 1 Choose **Management Appliance > Centralized Services > Web > Centralized Reporting**.

Step 2 Click **Edit Settings**.

Step 3 Select the **Anonymize usernames in reports** checkbox.

Step 4 Submit and commit your change.

Working with Interactive Web Reporting Pages

Interactive Web reporting pages allow you to monitor information on one or all of the managed Web Security appliances in your system.

To work with these pages, see the following topics:

Table 5-1 Working with Interactive Web Reporting Pages

To	See
View options for accessing and viewing report data	Ways to View Reporting Data, page 3-1
Understand the meaning of the data in tables	Table Column Descriptions for Web Reports, page 5-10
Customize your view of the interactive report pages	Customizing Your View of Report Data, page 3-3
Find information within your data	Web Tracking, page 5-50
Print or export report information	Printing and Exporting Reporting and Tracking Data, page 3-8
Understand the various interactive report pages	Understanding the Web Reporting Pages, page 5-7
Generate a report on demand	About Scheduled and On-Demand Web Reports, page 5-61
Schedule reports to run automatically at intervals and times that you specify	About Scheduled and On-Demand Web Reports, page 5-61
View archived on-demand and scheduled reports	Viewing and Managing Archived Web Reports, page 5-66
Understand how data is gathered	How the Security Appliance Gathers Data for Reports, page 3-2

Understanding the Web Reporting Pages

The **Web > Reporting** tab provides several options for viewing reporting data. This section describes each of the reporting pages under this tab, and explains the information displayed on each of the reporting pages.

Note For information on which of the options on the Web Reporting tab are available as on-demand or scheduled reports, see the "About Scheduled and On-Demand Web Reports" section on page 5-61.

Table 5-2 Web Reporting Tab Details

Web Reporting Menu	Action
Web Reporting Overview	The Overview page provides a synopsis of the activity on your Cisco IronPort appliances. It includes graphs and summary tables for the incoming and outgoing transactions. For more information, see the "Web Reporting Overview" section on page 5-12.
Users Report (Web)	The Users page provides several web tracking links that allows you to view web tracking information for individual users. From the **Users** page you can view how long a user, or users, on your system have spent on the internet, on a particular site or URL, and how much bandwidth that user is using. From the **Users** page you can click on an individual user in the interactive Users table to view more details for that specific user on the User Details page. The **User Details** page allows you to see specific information about a user that you have identified in the Users table on the **Web > Reporting > Users** page. From this page you can investigate individual user's activity on your system. This page is particularly useful if you are running user-level investigations and need to find out, for example, what sites your users are visiting, what Malware threats they are encountering, what URL categories they are accessing, and how much time a specific user is spending at these sites. For more information, see the "Users Report (Web)" section on page 5-16. For information on a specific user in your system, see the "User Details (Web Reporting)" section on page 5-19
Web Sites Report	The Web Sites page allows you to view an overall aggregation of the activity that is happening on your managed appliances. From this page you can monitor high-risk web sites accessed during a specific time range. For more information, see the "Web Sites Report" section on page 5-22.
URL Categories Report	The URL Categories page allows you to view the top URL Categories that are being visited, including: • the top URLs that have triggered a block or warning action to occur per transaction. • all the URL categories during a specified time range for both completed, warned and blocked transactions. This is an interactive table with interactive column headings that you can use to sort data as you need. For more information, see the "URL Categories Report" section on page 5-24.

Table 5-2 Web Reporting Tab Details

Web Reporting Menu	Action
Application Visibility Report	The Application Visibility page allows you to apply and view the controls that have been applied to a particular application types within the Security Management appliance and Web Security appliance. For more information, see the "Application Visibility Report" section on page 5-27.
Anti-Malware Report	The Anti-Malware page allows you to view information about malware ports and malware sites that the anti-malware scanning engine(s) detected during the specified time range. The upper part of the report displays the number of connections for each of the top malware ports and web sites. The lower part of the report displays malware ports and sites detected. For more information, see the "Anti-Malware Report" section on page 5-31.
Client Malware Risk Report	The Client Malware Risk page is a security-related reporting page that can be used to identify individual client computers that may be connecting unusually frequently to malware sites. For more information, see the "Client Malware Risk Report" section on page 5-37.
Web Reputation Filters Report	Allows you to view reporting on Web Reputation filtering for transactions during a specified time range. For more information, see the "Web Reputation Filters Report" section on page 5-39.
L4 Traffic Monitor Report	Allows you to view information about malware ports and malware sites that the L4 Traffic Monitor detected during the specified time range. For more information, see the "L4 Traffic Monitor Report" section on page 5-42.
SOCKS Proxy Report	Allows you to view data for SOCKS proxy transactions, including destinations and users. For more information, see the "SOCKS Proxy Report" section on page 5-46.
Reports by User Location	The Reports by User Location page allows you to find out what activities that your mobile users are conducting from their local or remote systems. For more information, see the "Reports by User Location" section on page 5-48.

Table 5-2 Web Reporting Tab Details

Web Reporting Menu	Action
Web Tracking	The Web Tracking page allows you to search for the following types of information: • Searching for Transactions Processed by Web Proxy Services allows you to track and see basic web-related information such as the type of web traffic that is being handled by the appliances. This includes information such as time ranges, and UserID and Client IP addresses, but also includes information like certain types of URLs, how much bandwidth that each connection is taking up, or tracking a specific user's web usage. • Searching for Transactions Processed by the L4 Traffic Monitor allows you to search your L4TM data for sites, ports, and client IP addresses involved in malware transfer activity. • Searching for Transactions Processed by the SOCKS Proxy allows you to search for transactions processed by the SOCKS proxy. For more information, see the "Web Tracking" section on page 5-50.
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" section on page 5-55.
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 "Data Availability Page" section on page 5-60.
Scheduled Reports	Allows you to schedule reports for a specified time range. For more information, see the "About Scheduled and On-Demand Web Reports" section on page 5-61.
Archived Reports	Allows you to archive reports for a specified time range. For more information, see the "Viewing and Managing Archived Web Reports" section on page 5-66.

Note You can schedule reports for most of the web reporting categories, including additional reports for Extended Top URL Categories and Top Application Types. For more information on scheduling reports, see the "About Scheduled and On-Demand Web Reports" section on page 5-61.

Table Column Descriptions for Web Reports

This section describes the column headings used in the tables on various Web report pages.

Note Not every column is available for every report page, and not every available column is visible by default. To view the columns available for a table, click the Column link below the table.

For more information about working with tables in reports, see Customizing Tables on Report Pages, page 3-6.

Table 5-3 Table Column Descriptions for Web Reporting Pages

Column Name	Description
Domain or Realm	The domain or realm of the user displayed in text format.
UserID or Client IP	The user ID or client IP of the user displayed in text format.
Bandwidth Used	The amount of bandwidth that is used by a particular user or action. Bandwidth units are displayed in Bytes or percentage.
Bandwidth Saved by Blocking	The amount of bandwidth that has been saved due to blocking certain transactions. Bandwidth units are displayed in Bytes
Time Spent	The amount of time spent on a web page. For purposes of investigating a user, the time spent by the user on each URL category. When tracking a URL, the time spent by each user on that specific URL. Once a transaction event is tagged as 'viewed', that is, a user goes to a particular URL, a 'Time Spent' value will start to be calculated and added as a field in the web reporting table. To calculate the time spent, AsyncOS assigns each active user with 60 seconds of time for activity during a minute. At the end of the minute, the time spent by each user is evenly distributed among the different domains the user visited. For example, if a user goes to four different domains in an active minute, the user is considered to have spent 15 seconds at each domain. For the purposes of the time spent value, considering the following notes: • An active user is defined as a user name or IP address that sends HTTP traffic through the appliance and has gone to a website that AsyncOS considers to be a "page view." • AsyncOS defines a page view as an HTTP request initiated by the user, as opposed to a request initiated by the client application. AsyncOS uses a heuristic algorithm to make a best effort guess to identify user page views. Units are displayed in Hours:Minutes format.

Table 5-3 Table Column Descriptions for Web Reporting Pages

Column Name	Description
Allowed URL Category	The number and type of categories that have been allowed. Units displayed in transaction type.
Monitored URL Category	The number and type of categories that are being monitored. Units displayed in transaction type.
Warned URL Category	The number and type of categories that have initiated a warning. Units displayed in transaction type.
Blocked by URL Category	The transaction that has been blocked due to URL Category. Units displayed in transaction type.
Blocked by Application or Application Type	The application that has been blocked due to application type. Units displayed in transaction type.
Blocked by Web Reputation	The transaction that has been blocked due to web reputation. Units displayed in transaction type.
Blocked by Anti-Malware	The transactions blocked by Anti-Malware. Units displayed in transaction type.
Other Blocked Transactions	All other transactions that have been blocked. Units displayed in transaction type.
Transactions with Bandwidth Limit	The number of transactions that have a bandwidth limit.
Transactions without Bandwidth Limit	The number of transactions that do not have a bandwidth limit.
Transactions Blocked by Application	The number of transactions blocked by a specific application type.
Warned Transactions	All transactions that rendered a warning to the user. Units displayed in transaction type.
Transactions Completed	The transactions completed by a user. Units displayed in transaction type.
Transactions Blocked	All transactions that have been blocked. Units displayed in transaction type.
Total Transactions	The total number of transactions that have occurred.

Web Reporting Overview

The **Web > Reporting > Overview** page provides a synopsis of the activity on your Cisco IronPort appliances. It includes graphs and summary tables for the incoming and outgoing transactions.

Figure 5-3 shows the Overview page.

Figure 5-3 The Web > Reporting > Overview Page

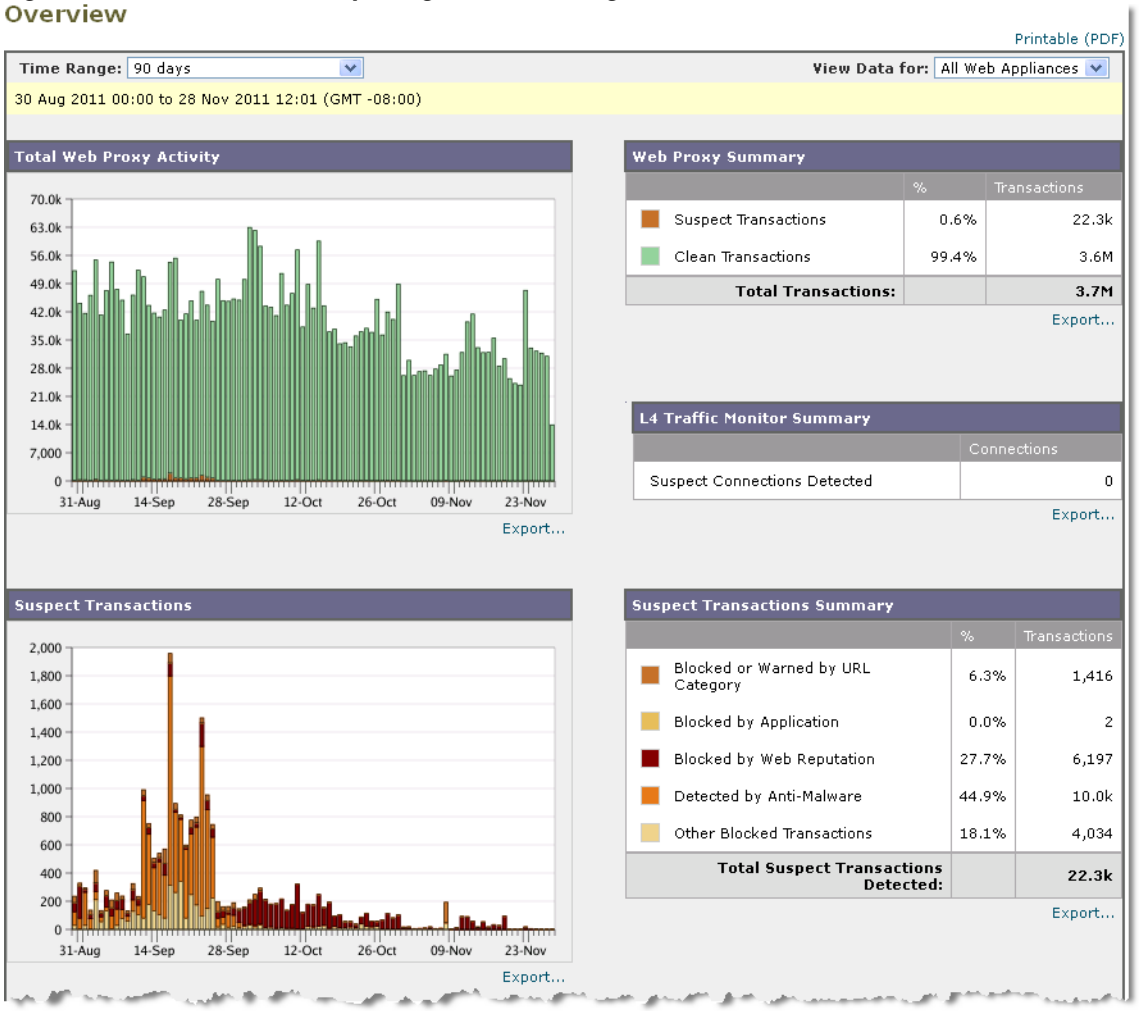

(Illustration continues on next page)

(Illustration continued from previous page)

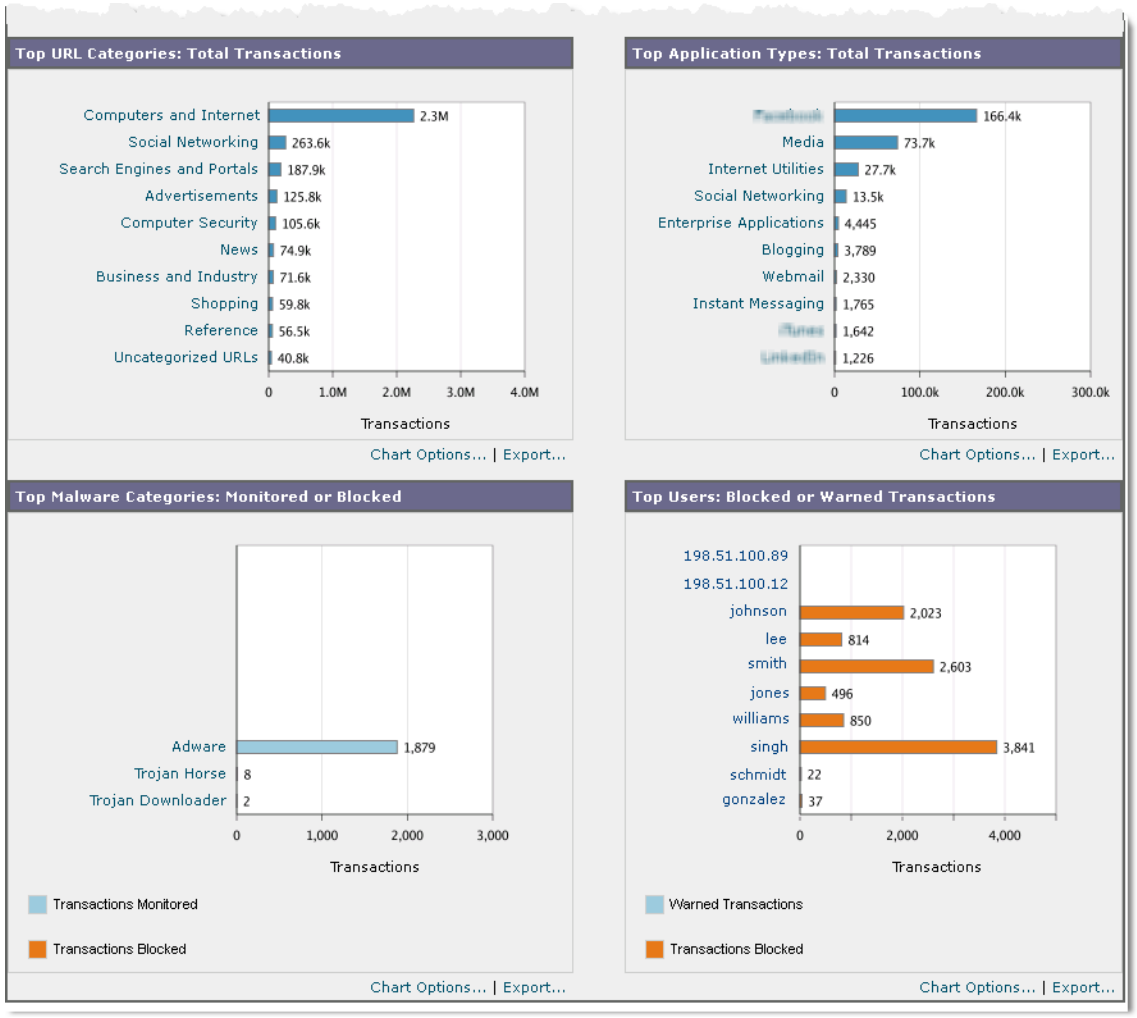

At a high level the **Overview** page shows you statistics about the URL and User usage, Web Proxy activity, and various transaction summaries. The transaction summaries gives you further trending details on, for example suspect transactions, and right across from this graph, how many of those suspect transactions are blocked and in what manner they are being blocked.

The lower half of the **Overview** page is about usage. That is, the top URL categories being viewed, the top application types and categories that are being blocked, and the top users that are generating these blocks or warnings.

The following list explains the various sections on the **Overview** page:

Table 5-4 Details on the Web > Reporting > Overview 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" section on page 3-4.
View Data for	Choose a Web Security appliance for which you want to view Overview data, or choose All Web Appliances. See also Viewing Reporting Data for an Appliance or Reporting Group, page 3-4.
Total Web Proxy Activity	This section allows you to view the web proxy activity that is being reported by the Web Security appliances that are currently managed by the Security Management appliance. This section displays the actual number of transactions (vertical scale) as well as the approximate date that the activity occurred (horizontal timeline).
Web Proxy Summary	This section allows you to view the percentage of web proxy activity that are suspect, or clean proxy activity, including the total number of transactions.
L4 Traffic Monitor Summary	This section reports any Layer 4 traffic that is being reported by the Web Security appliances that are currently managed by the Security Management appliance.
Suspect Transactions	This section allows you to view the web transactions that have been labeled as suspect by the administrator. This section displays the actual number of transactions (vertical scale) as well as the approximate date that the activity occurred (horizontal timeline).
Suspect Transactions Summary	This section allows you to view the percentage of blocked or warned transactions that are suspect. Additionally you can see the type of transactions that have been detected and blocked, and the actual number of times that this transaction was blocked.
Top URL Categories by Total Transactions	This section displays the top 10 URL categories that are being blocked, including the type of URL category (vertical scale) and the actual number of times the specific type of category has been blocked (horizontal scale). The set of predefined URL categories is occasionally updated. For more information about the impact of these updates on report results, see URL Category Set Updates and Reports, page 5-26.
Top Application Types by Total Transactions	This section displays the top application types that are being blocked, including the name of the actual application type (vertical scale) and the number of times the specific application has been blocked (horizontal scale).

Table 5-4 Details on the Web > Reporting > Overview Page

Section	Description
Top Malware Categories Detected	This section displays all Malware categories that have been detected.
Top Users Blocked or Warned Transactions	This section displays the actual users that are generating the blocked or warned transactions. Users can be displayed by IP address or by user name. To make user names unrecognizable, see Anonymizing User Names in Web Reports, page 5-4.

Users Report (Web)

The **Web > Reporting > Users** page provides several links that allow you to view web reporting information for individual users.

From the **Users** page you can view how long a user, or users, on your system have spent on the internet, on a particular site or URL, and how much bandwidth that user is using.

Note The maximum number of users on the Web Security appliance that the Security Management appliance can support is 500.

Figure 5-4 Web > Reporting > Users Page

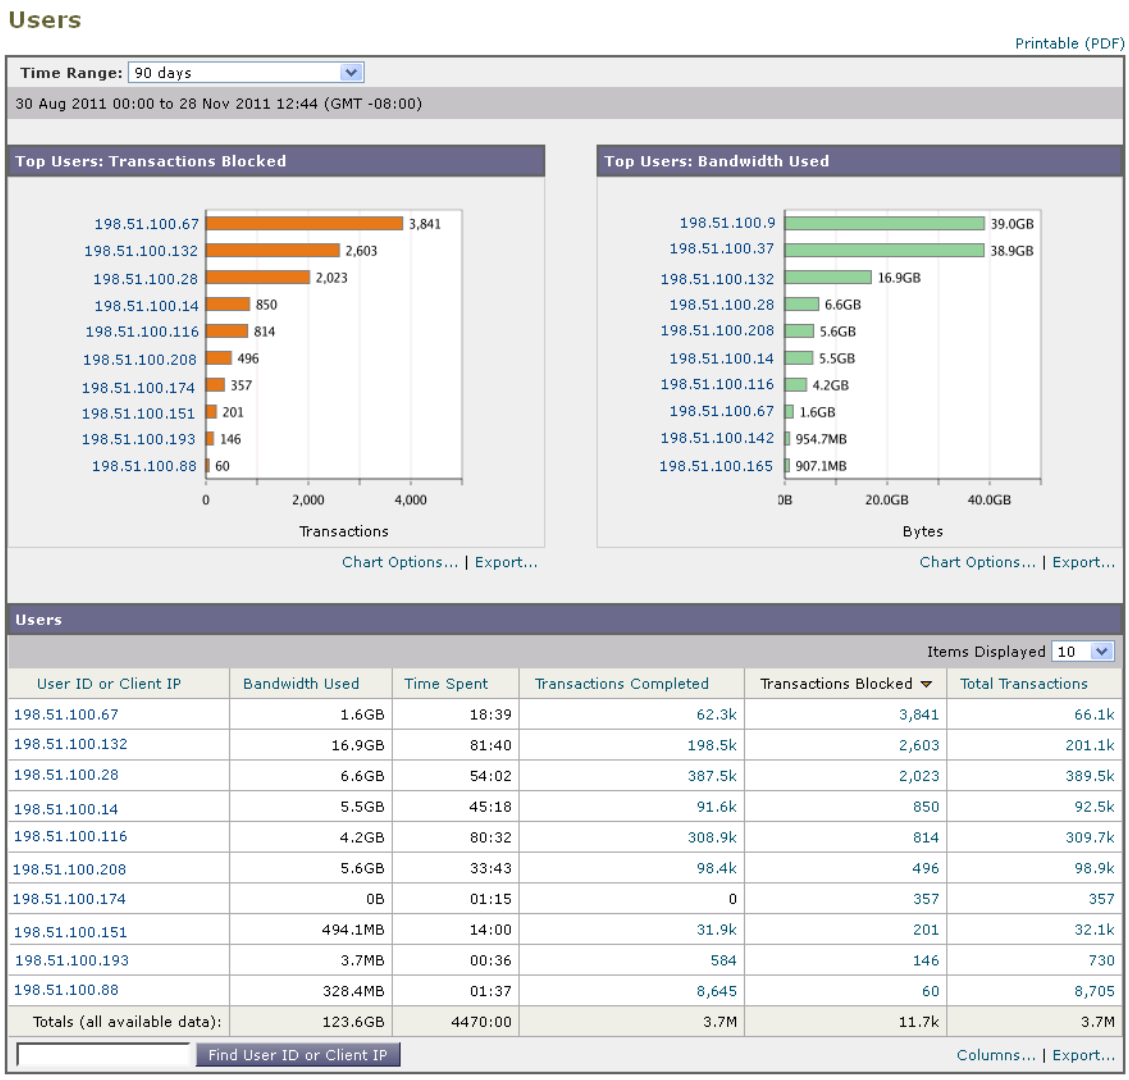

From the **Users** page, you can view the following information pertaining to the users on your system:

Table 5-5 Details on the Web > Reporting > 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" section on page 3-4.
Top Users by Transactions Blocked	This section lists the top users, by either IP address or user name (vertical scale), and the number of transactions that have been blocked specific to that user (horizontal scale). The user name or IP address can be made unrecognizable for reporting purposes. For more information on how to make user names unrecognizable in for this page or in scheduled reports, see the section "Enabling Centralized Web Reporting on the Security Management Appliance" section on page 5-3. The default setting is that all user names appear. To hide user names, see "Anonymizing User Names in Web Reports" section on page 5-4.
Top Users by Bandwidth Used	This sections displays the top users, by either IP address or user name (vertical scale), that are using the most bandwidth on the system (horizontal scale represented in gigabyte usage).
Users Table	For more information about the data in this table, see Table Column Descriptions for Web Reports, page 5-10. Additionally, you can find a specific User ID or Client IP address. In the text field at the bottom of the User section, enter the specific User ID or Client IP address and click on **Find User ID or Client IP Address**. The IP address does not need to be an exact match to return results. From the Users table you can click on a specific user to find more specific information. This information appears on the User Details page. For more information on the User Details page, see the "User Details (Web Reporting)" section on page 5-19

Note To view user IDs instead of client IP addresses, you must set up your Security Management appliance to obtain user information from an LDAP server. For information, see Creating the LDAP Server Profile in Chapter 9.

Tip To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

To view an example of how the **Users** page may be used, see "Example 1: Investigating a User" section on page D-1.

Note You can generate or schedule a report for the Users page. For information, see the "About Scheduled and On-Demand Web Reports" section on page 5-61.

User Details (Web Reporting)

The **User Details** page allows you to see specific information about a user that you have identified in the interactive Users table on the **Web > Reporting > Users** page.

The **User Details** page allows you to investigate individual user's activity on your system. This page is particularly useful if you are running user-level investigations and need to find out, for example, what sites your users are visiting, what Malware threats they are encountering, what URL categories they are accessing, and how much time a specific user is spending at these sites.

To display the **User Details** page for a specific user, click on a specific user from the User table on the **Web > Users** page and the following page appears:

Figure 5-5 User Details Page

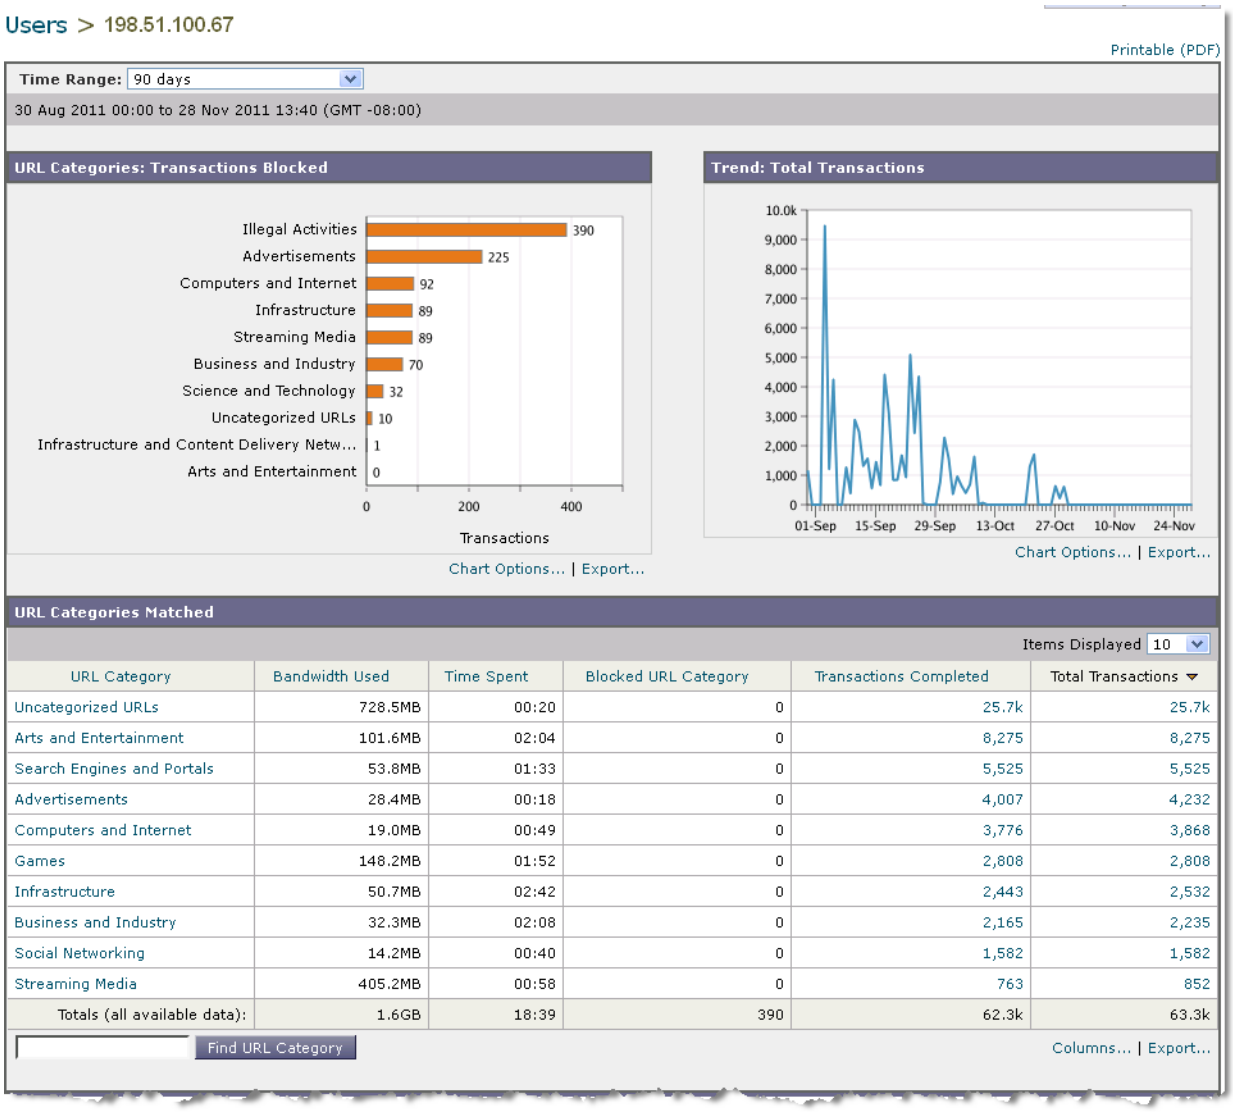

(Illustration continues on next page)

(Illustration continued from previous page)

Domains Matched

Items Displayed 10

Domain or IP	Bandwidth Used	Time Spent	Transactions Completed	Transactions Blocked	Total Transactions
[redacted].com	713.4MB	00:29	24.5k	0	24.5k
[redacted].com	92.4MB	02:02	8,037	0	8,037
[redacted].com	31.2MB	00:38	3,095	0	3,095
[redacted].com	1.7MB	00:56	179	1,769	1,948
[redacted].com	3.7MB	00:00	1,841	0	1,841
[redacted].com	2.4MB	02:12	1,539	80	1,619
[redacted].tv	12.6MB	00:03	1,033	0	1,033
[redacted].net	10.5MB	00:00	1,001	0	1,001
[redacted].com	46.5KB	00:57	4	898	902
[redacted].com	8.8MB	00:09	778	0	778

Find Domain or IP | Columns... | Export...

Applications Matched

Items Displayed 10

Application	Application Type	Bandwidth Used	Transactions Completed	Other Blocked Transactions	Total Transactions
[redacted]	Internet Utilities	3.7MB	1,832	0	1,832
[redacted] Video	Media	380.6MB	1,517	0	1,517
[redacted] General	[redacted]	10.2MB	1,283	0	1,283
[redacted]	Media	274.2MB	517	0	517
[redacted]	Instant Messaging	337.3KB	95	0	95
[redacted]	Webmail	1.4MB	68	0	68
[redacted] Mail	Webmail	425.8KB	61	0	61
[redacted]	Social Networking	364.5KB	58	0	58
[redacted] Photos	[redacted]	2.2MB	54	0	54
[redacted]	Media	157.6MB	40	0	40
Totals (all available data):	--	832.1MB	5,621	0	5,621

Find Application | Columns... | Export...

Malware Threats Detected

Malware Threat	Malware Category	Bandwidth Saved by Blocking	Transactions Monitored	Transactions Blocked	Total Malware Transactions Detected
[redacted]	Adware	0B	82	0	82
[redacted]	Adware	0B	8	0	8
[redacted]	Trojan Horse	36.0KB	0	3	3
Totals (all available data):	--	36.0KB	90	3	93

Find Malware Threat | Columns... | Export...

Policies Matched

Policy Name	Policy Type	Bandwidth Used	Completed Transactions	Blocked Transactions	Total Transactions
Policy 1	Access	1.6GB	62.2k	1,174	63.4k
Policy 2	Access	0B	0	2,667	2,667
Policy 3	Decryption	768.3KB	91	0	91
Totals (all available data):	--	1.6GB	62.3k	3,841	66.1k

Find Policy Name | Columns... | Export...

From the **User Details** page, you can view the following information pertaining to an individual user on your system:

Table 5-6 Details on the Web > Reporting > User > User Details Page

Section	Description
Time Range (drop-down list)	A menu that allows you to choose the time range of the data contained in the report. For more information on time ranges and customizing this for your needs, see the "Choosing a Time Range for Reports" section on page 3-4.
URL Categories by Total Transactions	This section lists the specific URL Categories that a specific user is using. The set of predefined URL categories is occasionally updated. For more information about the impact of these updates on report results, see URL Category Set Updates and Reports, page 5-26.
Trend by Total Transactions	This graph displays at what times the user accessed the web. For example, this graph will indicate if there is a large spike in web traffic during certain hours of the day, and when those spikes occur. Using the Time Range drop-down list, you can expand this graph to see a more or less granular span of time that this user was on the web.
URL Categories Matched	The URL Categories Matched section shows matched categories for both completed and blocked transactions. From this section you can also find a specific URL Category. In the text field at the bottom of the section enter the URL Category and click **Find URL Category**. The category does not need to be an exact match. The set of predefined URL categories is occasionally updated. For more information about the impact of these updates on report results, see URL Category Set Updates and Reports, page 5-26.
Domains Matched	From this section you can find out about a specific Domain or IP address that this user has accessed. You can also see the time spent on those categories, and various other information that you have set from the column view. In the text field at the bottom of the section enter the Domain or IP address and click **Find Domain or IP**. The domain or IP address does not need to be an exact match.
Applications Matched	From this section you can find a specific application that a specific user is using. For example, if a user is accessing a site that requires use of a lot of Flash video, you will see the application type in the Application column. In the text field at the bottom of the section enter the application name and click **Find Application**. The name of the application does not need to be an exact match.

Table 5-6 Details on the Web > Reporting > User > User Details Page

Section	Description
Malware Threats Detected	From this table you can see the top Malware threats that a specific user is triggering. You can search for data on a specific malware threat name in the Find Malware Threat field. Enter the Malware Threat name and click Find Malware Threat. The name of the Malware Threat does not need to be an exact match.
Policies Matched	From this section you can find the policy groups that applied to this user when accessing the web. In the text field at the bottom of the section enter the policy name and click **Find Policy**. The name of the policy does not need to be an exact match.

Note From Client Malware Risk Details table: The client reports sometimes show a user with an asterisk (*) at the end of the user name. For example, the Client report might show an entry for both "jsmith" and "jsmith*". User names listed with an asterisk (*) indicate the user name provided by the user, but not confirmed by the authentication server. This happens when the authentication server was not available at the time and the appliance is configured to permit traffic when authentication service is unavailable.

To view an example of how the User Details page may be used, see "Example 1: Investigating a User" section on page D-1.

Web Sites Report

The **Web > Reporting > Web Sites** page is an overall aggregation of the activity that is happening on the managed appliances. From this page you can monitor high-risk web sites accessed during a specific time range.

Figure 5-6 Web Sites Page

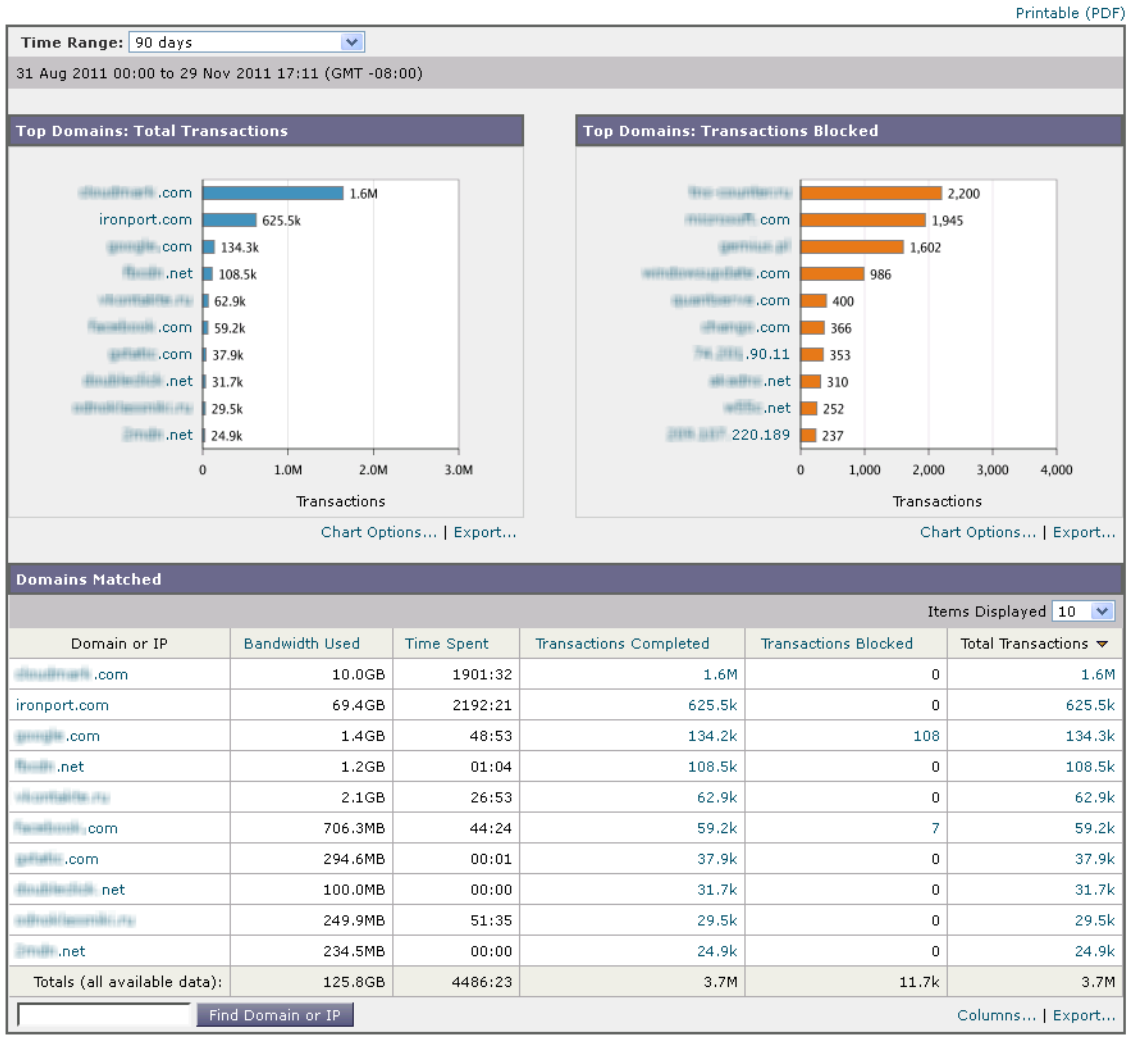

From the **Web Sites** page, you can view the following information:

Table 5-7 Details on the Web > Reporting > Web Sites 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" section on page 3-4.
Top Domains by Total Transactions	This section lists the top domains that are being visited on the site in a graph format.

Table 5-7 Details on the Web > Reporting > Web Sites Page

Section	Description
Top Domains by Transactions Blocked	This section lists the top domains that triggered a block action to occur per transaction in a graph format. For example, a user went to a certain domain and because of a specific policy that I have in place, this triggered a block action. This domain is listed in this graph as a transaction blocked, and the domain site that triggered the block action is listed.
Domains Matched	This section lists the domains that are that are being visited on the site in an interactive table. From this table you can access more granular information about a specific domain by clicking on the specific domain. The Proxy Services tab on the Web Tracking page appears and you can see tracking information and why certain domains were blocked. For more information about the data in this table, see Table Column Descriptions for Web Reports, page 5-10. When you click on a specific domain you can see the top users of that domain, the top transactions on that domain, the URL Categories matched and the Malware threats that have been detected. To view an example of how Web Tracking may be used, see "Example 2: Tracking a URL" section on page D-5.

Tip To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

Note You can generate or schedule a report for information on the Web Sites page. For information, see the "About Scheduled and On-Demand Web Reports" section on page 5-61.

URL Categories Report

The **Web > Reporting > URL Categories** page can be used to view the URL categories of sites that users on your system are visiting.

Figure 5-7 URL Categories Page

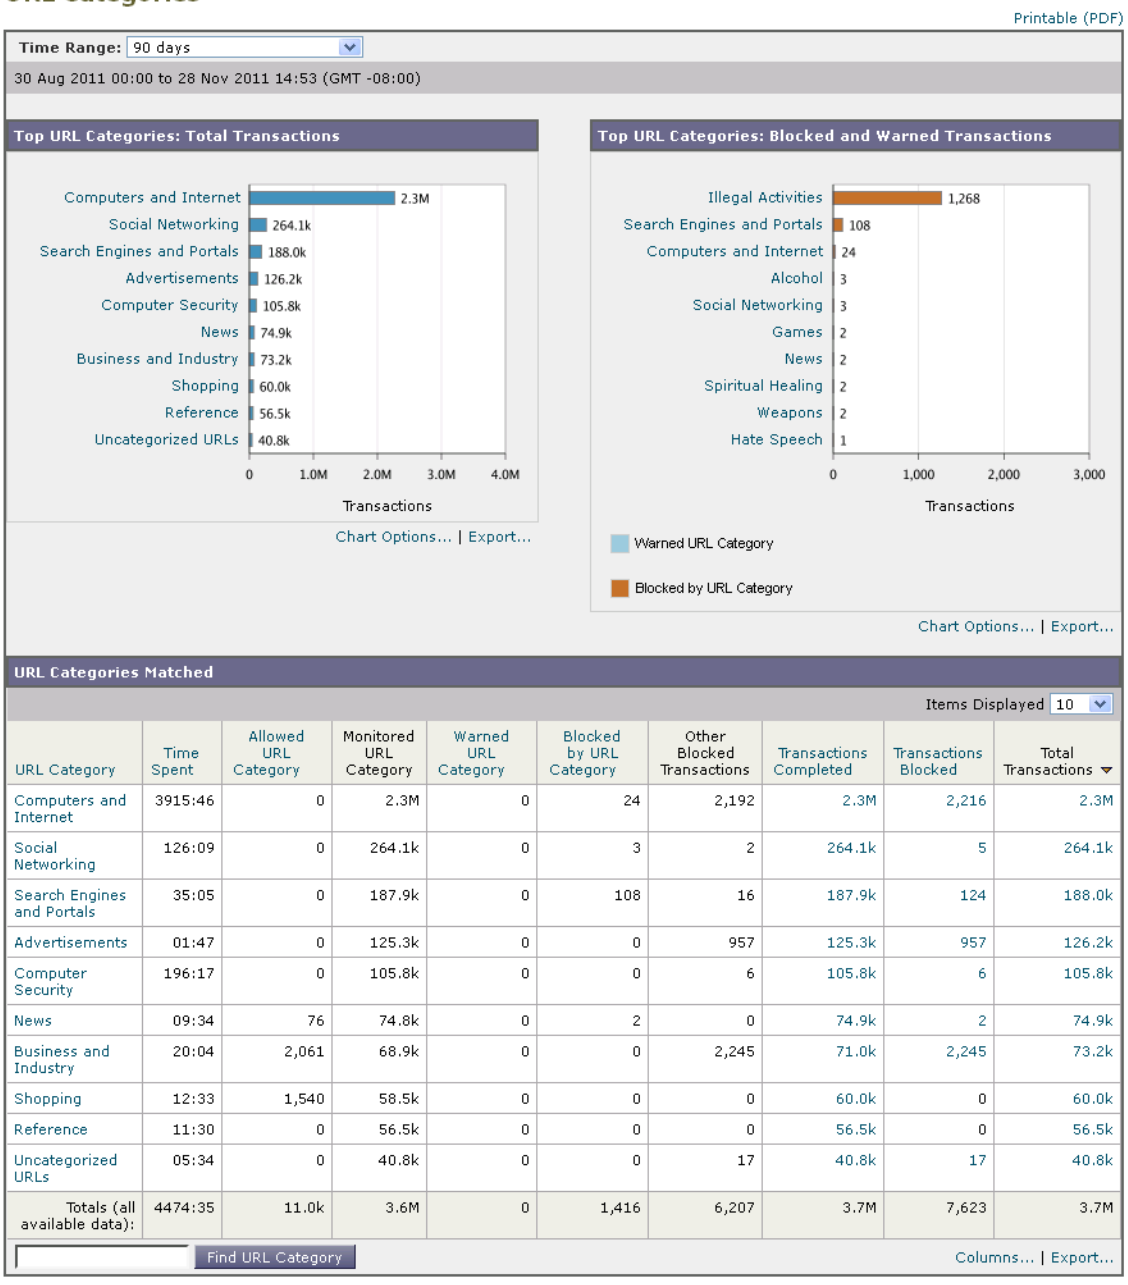

From the **URL Categories** page, you can view the following information:

Table 5-8 Details on the Web > Reporting > URL Categories Page

Section	Description
Time Range (drop-down list)	Choose the time range for your report. For more information, see the "Choosing a Time Range for Reports" section on page 3-4.
Top URL Categories by Total Transactions	This section lists the top URL Categories that are being visited on the site in a graph format.

Table 5-8 Details on the Web > Reporting > URL Categories Page

Section	Description
Top URL Categories by Blocked and Warned Transactions	This section lists the top URL that triggered a block or warning action to occur per transaction in a graph format. For example, a user went to a certain URL and because of a specific policy that is in place, this triggered a block action or a warning. This URL then gets listed in this graph as a transaction blocked or warning.
URL Categories Matched	The URL Categories Matched section shows the disposition of transactions by URL category during the specified time range, plus bandwidth used and time spent in each category. If the percentage of uncategorized URLs is higher than 15-20%, consider the following options: • For specific localized URLs, you can create custom URL categories and apply them to specific users or group policies. For more information, see the "Custom URL Categories" section of the *Cisco IronPort AsyncOS for Web Security User Guide*. • For sites that you feel should be included in existing or other categories, see Reporting Misclassified and Uncategorized URLs, page 5-27.

Tip To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

Note
- To generate a more detailed report than this page can provide, see Top URL Categories—Extended, page 5-63.
- If Data Availability is used within a scheduled report for URL Categories, and there are gaps in data for any of the appliances, the following message is displayed at the bottom of the page: "`Some data in this time range was unavailable.`" If there are no gaps present, nothing appears.

URL Category Set Updates and Reports

The set of predefined URL categories may periodically be updated on your Security Management appliance, as described in URL Category Set Updates and Centralized Configuration Management, page 8-22.

When these updates occur, data for old categories will continue to appear in reports and web tracking results until the data is too old to be included. Report data generated after a category set update will use the new categories, so you may see both old and new categories in the same report.

If there is overlap between the contents of old and new categories, you may need to examine report results more carefully to obtain valid statistics. For example, if the "Instant Messaging" and "Web-based Chat" categories have been merged into a single "Chat and Instant Messaging" category during the time frame that you are looking at, visits before the merge to sites covered by the "Instant Messaging" and

"Web-based Chat" categories are not counted in the total for "Chat and Instant Messaging". Likewise, visits to instant messaging or Web-based chat sites after the merge would not be included in the totals for the "Instant Messaging" or "Web-based Chat" categories.

Using The URL Categories Page in Conjunction with Other Reporting Pages

The URL Categories page can be used in conjunction with the Application Visibility Report and the Users Report (Web) to investigate a particular user and the types of applications or websites that a particular user is trying to access.

For example, from the URL Categories Report you can generate a high level report for Human Resources which details all the URL categories that are visited by the site. From the same page, you can gather further details in the URL Categories interactive table about the URL category 'Streaming Media'. By clicking on the Streaming Media category link, you can view the specific URL Categories report page. This page not only displays the top users that are visiting streaming media sites (in the Top Users by Category for Total Transactions section), but also displays the domains that are visited (in the Domains Matched interactive table) such as YouTube.com or QuickPlay.com.

At this point, you are getting more and more granular information for a particular user. Now, let's say this particular user stands out because of their usage, and you want to find out exactly what they are accessing. From here you can click on the user in the Users interactive table. This action takes you to the User Details (Web Reporting), where you can view the user trends for that user, and find out exactly what they have been doing on the web.

If you wanted to go further, you can now get down to web tracking details by clicking on Transactions Completed link in the interactive table. This displays the Searching for Transactions Processed by Web Proxy Services on the Web Tracking page where you can see the actual details about what dates the user accessed the sites, the full URL, the time spent on that URL, etc.

To view another example of how the URL Categories page may be used, see "Example 3: Investigating Top URL Categories Visited" section on page D-6.

Reporting Misclassified and Uncategorized URLs

You can report misclassified and uncategorized URLs at the following URL:

`https://securityhub.cisco.com/web/submit_urls`

Submissions are evaluated for subsequent rule updates.

To check the status of submitted URLs, click the **Status on Submitted URLs** tab on this page.

Application Visibility Report

Note For detailed information on Application Visibility, see the 'Understanding Application Visibility and Control' chapter in the *Cisco IronPort AsyncOS for Web Security User Guide.*

The **Web > Reporting > Application Visibilit**y page allows you to apply controls to particular application types within the Security Management appliance and Web Security appliance.

Not only does application control gives you more granular control over web traffic than just URL filtering, for example, it gives you more control over the following types of applications, and application types:

- Evasive applications, such as anonymizers and encrypted tunnels.
- Collaboration applications, such as Cisco WebEx, Facebook, and instant messaging.
- Resource intensive applications, such as streaming media.

Understanding the Difference between Application versus Application Types

It is crucial to understand the difference between an application and an application types so that you can control the applications involved for your reports.

- **Application Types.** A category that contains one or more applications. For example, **search engines** is an application type that may contain search engines such as Google Search and Craigslist. Instant messaging is another application type category which may contain Yahoo Instant Messenger, or Cisco WebEx. Facebook is also an application type.
- **Applications.** Particular applications that belong in an application type. For example, YouTube is an application in the Media application type.
- **Application behaviors.** Particular actions or behaviors that users can accomplish within an application. For example, users can transfer files while using an application, such as Yahoo Messenger. Not all applications include application behaviors you can configure.

Note For detailed information on understanding how you can use Application Visibility and Control (AVC) engine to control Facebook activity, see the ‘Understanding Application Visibility and Control’ chapter in the *Cisco IronPort AsyncOS for Web Security User Guide.*

Figure 5-8 Application Visibility Page

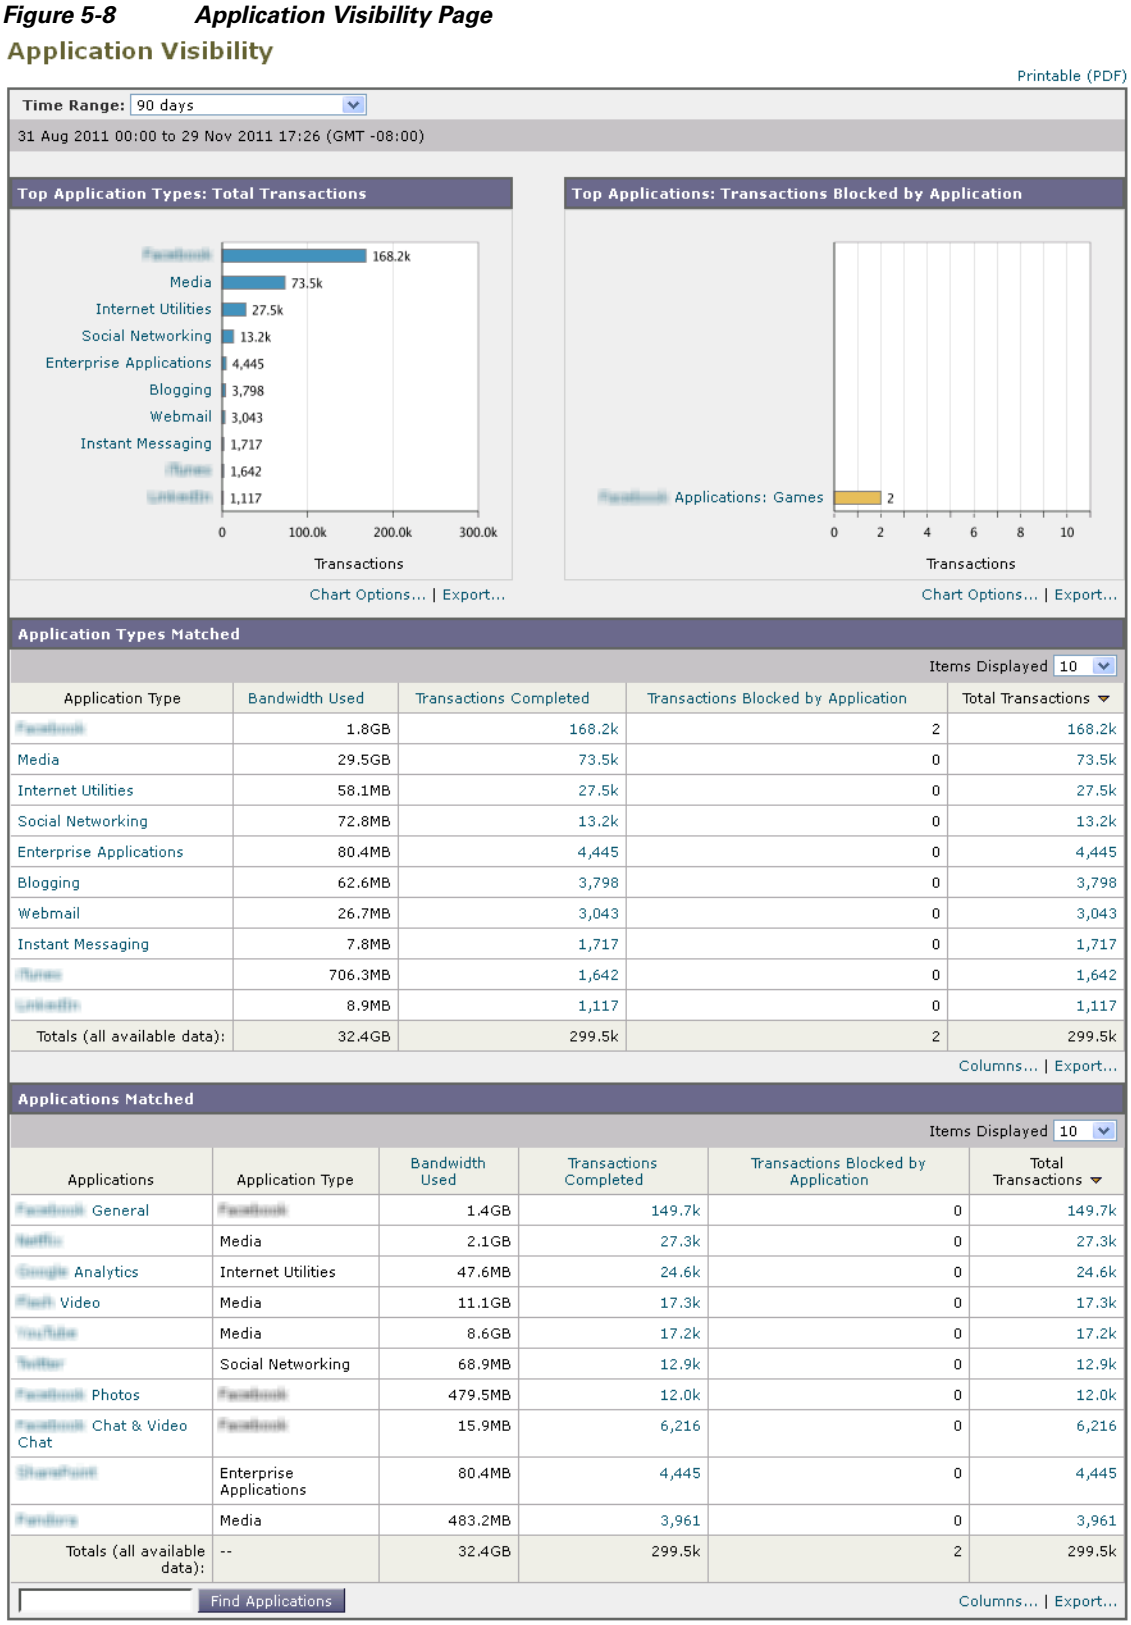

From the **Application Visibility** page, you can view the following information:

Table 5-9 Details on the Web > Reporting > Application Visibility 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" section on page 3-4.
Top Application Types by Total Transactions	This section lists the top application types that are being visited on the site in a graph format. For example, instant messaging tools such as Yahoo Instant Messenger, Facebook, and Presentation application types.
Top Applications by Blocked Transactions	This section lists the top application types that triggered a block action to occur per transaction in a graph format. For example, a user has tried to start a certain application type, for example Google Talk or Yahoo Instant Messenger, and because of a specific policy that is in place, this triggered a block action. This application then gets listed in this graph as a transaction blocked or warning.
Application Types Matched	The Application Types Matched interactive table allows you to view granular details about the application types listed in the Top Applications Type by Total Transactions table. From the Applications column you can click on an application to view details
Applications Matched	The Applications Matched section shows all the application during a specified time range. This is an interactive table with interactive column headings that you can use to sort data as you need. You can configure the columns that you want to appear in the Applications Matched section. For information on configuring columns for this section, see the "Working with Interactive Web Reporting Pages" section on page 5-7. After you have selected the specific items to appear in the Applications table, you can select how many items you want to be displayed from the **Items Displayed** drop-down menu. Choices are: 10, 20, 50, or 100. Additionally, you can find a specific Application within the Application Matched section. In the text field at the bottom of this section, enter the specific Application name and click **Find Application**.

Tip To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

Note You can generate a scheduled report for information on the Application Visibility page. For information on scheduling a report, see the "About Scheduled and On-Demand Web Reports" section on page 5-61.

Anti-Malware Report

The **Web > Reporting> Anti-Malware** page is a security-related reporting page that reflects the results of scanning by your enabled scanning engines (Webroot, Sophos, McAfee, and/or Adaptive Scanning).

Use this page to help identify and monitor web-based malware threats.

Note To view data for malware found by L4 Traffic Monitoring, see "L4 Traffic Monitor Report" section on page 5-42.

Figure 5-9 Anti-Malware Page

Anti-Malware

Printable (PDF)

Time Range: 90 days

31 Aug 2011 00:00 to 29 Nov 2011 17:37 (GMT -08:00)

Top Malware Categories: Monitored or Blocked

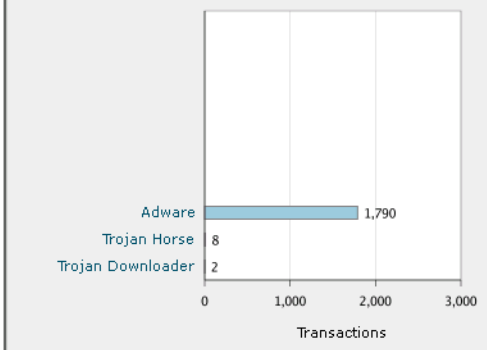

Chart Options... | Export...

Top Malware Threats: Monitored or Blocked

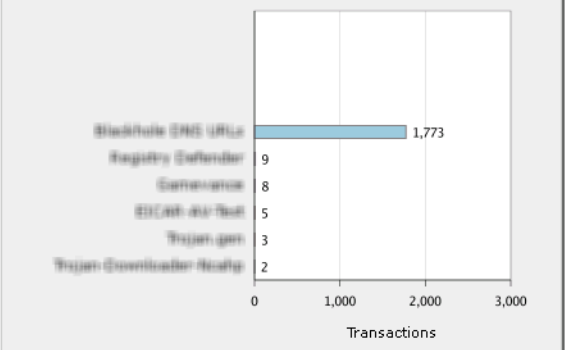

Chart Options... | Export...

Malware Categories

Malware Category	Bandwidth Saved by Blocking	Requests Monitored	Responses Monitored	Requests Blocked	Responses Blocked	Transactions Monitored	Transactions Blocked	Transactions Detected
Adware	0B	0	1,790	0	0	1,790	0	1,790
Trojan Horse	96.0KB	0	0	0	8	0	8	8
Trojan Downloader	0B	0	2	0	0	2	0	2
Totals (all available data):	96.0KB	0	1,792	0	8	1,792	8	1,800

Columns... | Export...

Malware Threats

Malware Threat	Malware Category	Bandwidth Saved by Blocking	Transactions Monitored	Transactions Blocked	Transactions Detected
[redacted]	Adware	0B	1,773	0	1,773
[redacted]	Adware	0B	9	0	9
[redacted]	Adware	0B	8	0	8
[redacted]	Trojan Horse	60.0KB	0	5	5
[redacted]	Trojan Horse	36.0KB	0	3	3
[redacted]	Trojan Downloader	0B	2	0	2
Totals (all available data):	--	96.0KB	1,792	8	1,800

Columns... | Export...

From the **Anti-Malware** page, you can view the following information:

Table 5-10 Details on the Web > Reporting > Anti-Malware 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" section on page 3-4.
Top Malware Categories: Monitored or Blocked	This section displays the top malware categories that are detected by a given category type. This information is displayed in graph format. See Malware Category Descriptions, page 5-35 for more information on valid Malware categories.
Top Malware Threats: Monitored or Blocked	This section displays the top malware threats. This information is displayed in graph format.
Malware Categories	The Malware Categories interactive table shows detailed information about particular malware categories that are displayed in the Top Malware Categories chart. Clicking on any of the links in the Malware Categories interactive table allows you to view more granular details about individual malware categories and where they are on the network. Exception: an Outbreak Heuristics link in the table lets you view a chart showing when transactions in this category occurred. See Malware Category Descriptions, page 5-35 for more information on valid Malware categories.
Malware Threats	The Malware Threats interactive table shows detailed information about particular malware threats that are displayed in the Top Malware Threats section. Threats labeled "Outbreak" with a number are threats identified by the Adaptive Scanning feature independently of other scanning engines. **Note** When sorting the table by Malware Threat in ascending order, Unnamed Malware appears at the top of the list.

Tip To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

Malware Category Report

The Malware Category Report page allows you to view detailed information on an individual Malware Category and what it is doing on your network.

To access the Malware Category report page, perform the following:

Procedure

Step 1 On the Security Management appliance window, choose **Web > Reporting > Anti-Malware**.

Step 2 In the Malware Categories interactive table, click on a category in the Malware Category column.

The Malware Category report page appears.

Figure 5-10 Malware Category Report Page

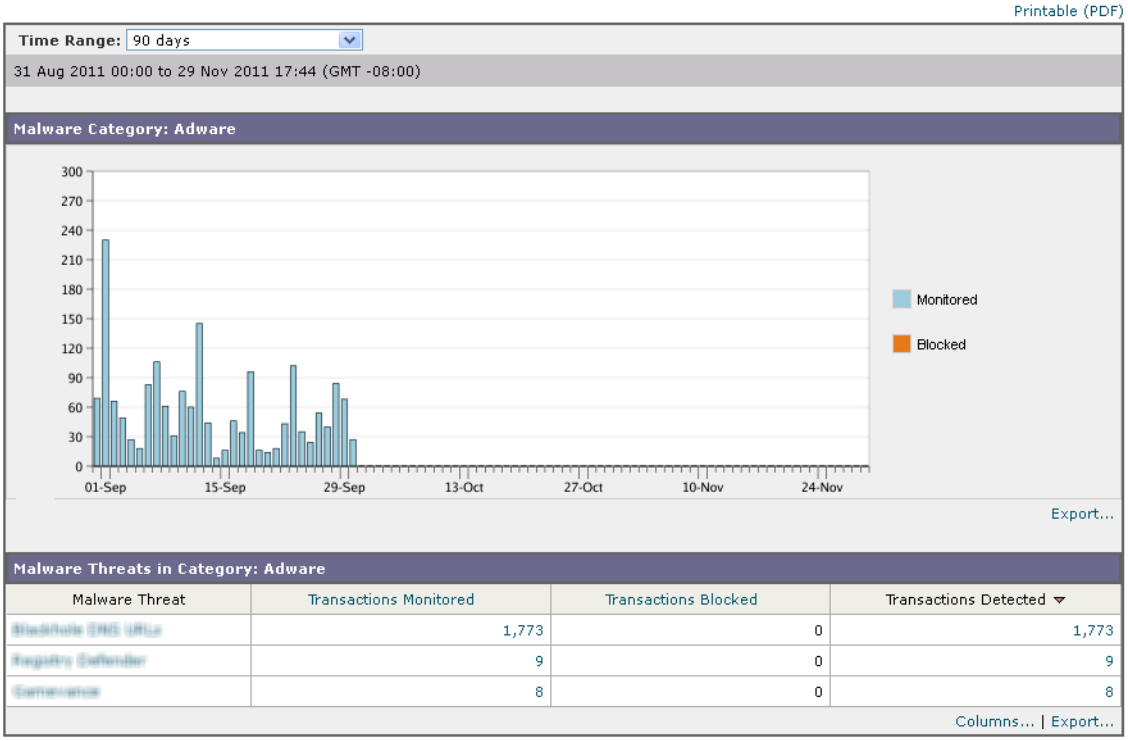

Step 3 To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

Malware Threat Report

The Malware Threat Report page report shows clients at risk for a particular threat, displays a list of potentially infected clients, and links to the Client Detail page. The trend graph at the top of the report shows monitored and blocked transactions for a threat during the specified time range. The table at the bottom shows the actual number of monitored and blocked transactions for a threat during the specified time range.

To access the Malware Threat report page, perform the following:

Procedure

Step 1 On the Security Management appliance window, choose **Web > Reporting > Anti-Malware**.

Step 2 In the Malware Threat interactive table, click on a category in the Malware Category column.

The Malware Threat report page appears:

Figure 5-11 Malware Threats Report Page

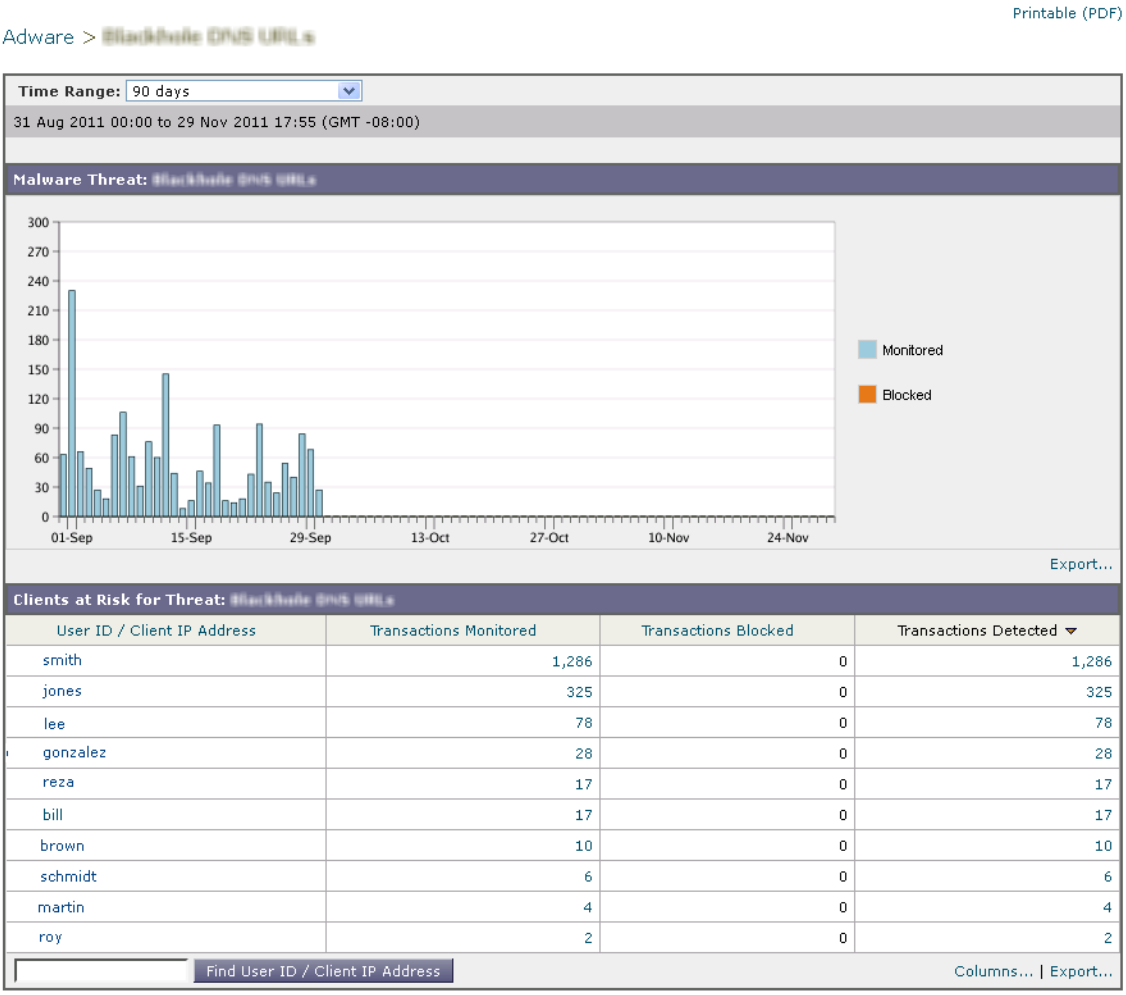

Step 3 To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

Step 4 For additional information, click the **Support Portal Malware Details** link below the table.

Note You can generate a scheduled report for Top Malware Categories Detected and Top Malware Threats Detected on the Anti-Malware page, but you cannot schedule a generated report from the Malware Categories and Malware Threats Report Page. For information on scheduling a report, see the "About Scheduled and On-Demand Web Reports" section on page 5-61.

Malware Category Descriptions

The Web Security appliance can block the following types of malware:

Malware Type	Description
Adware	Adware encompasses all software executables and plug-ins that direct users towards products for sale. Some adware applications have separate processes that run concurrently and monitor each other, ensuring that the modifications are permanent. Some variants enable themselves to run each time the machine is started. These programs may also change security settings making it impossible for users to make changes to their browser search options, desktop, and other system settings.
Browser Helper Object	A browser helper object is browser plug-in that may perform a variety of functions related to serving advertisements or hijacking user settings.
Commercial System Monitor	A commercial system monitor is a piece of software with system monitor characteristics that can be obtained with a legitimate license through legal means.
Dialer	A dialer is a program that utilizes your modem or another type of Internet access to connect you to a phone line or a site that causes you to accrue long distance charges to which you did not provide your full, meaningful, and informed consent.
Generic Spyware	Spyware is a type of malware installed on computers that collects small pieces of information about users without their knowledge.
Hijacker	A hijacker modifies system settings or any unwanted changes to a user's system that may direct them to a website or run a program without a user's full, meaningful, and informed consent.
Other Malware	This category is used to catch all other malware and suspicious behavior that does not exactly fit in one of the other defined categories.
Outbreak Heuristics	This category represents malware found by Adaptive Scanning independently of the other anti-malware engines.
Phishing URL	A phishing URL is displayed in the browser address bar. In some cases, it involves the use of domain names and resembles those of legitimate domains. Phishing is a form of online identity theft that employs both social engineering and technical subterfuge to steal personal identity data and financial account credentials.
PUA	Potentially Unwanted Application. A PUA is an application that is not malicious, but which may be considered to be undesirable.
System Monitor	A system monitor encompasses any software that performs one of the following actions: Overtly or covertly records system processes and/or user action. Makes those records available for retrieval and review at a later time.
Trojan Downloader	A trojan downloader is a Trojan that, after installation, contacts a remote host/site and installs packages or affiliates from the remote host. These installations usually occur without the user's knowledge. Additionally, a Trojan Downloader's payload may differ from installation to installation since it obtains downloading instructions from the remote host/site.
Trojan Horse	A trojan horse is a destructive program that masquerades as a benign application. Unlike viruses, Trojan horses do not replicate themselves.

Malware Type	Description
Trojan Phisher	A trojan phisher may sit on an infected computer waiting for a specific web page to be visited or may scan the infected machine looking for user names and passwords for bank sites, auction sites, or online payment sites.
Virus	A virus is a program or piece of code that is loaded onto your computer without your knowledge and runs against your wishes.
Worm	A worm is program or algorithm that replicates itself over a computer network and usually performs malicious actions.

Client Malware Risk Report

The **Web > Reporting > Client Malware Risk** page is a security-related reporting page that can be used to monitor client malware risk activity.

From the Client Malware Risk page, a system administrator can see which of their users are encountering the most blocks or warnings. Given the information gathered from this page, the administrator can click on the user link to view what this user doing on the web that makes them run into so many blocks or warnings and setting off more detections than the rest of the users on the network.

Additionally, the Client Malware Risk page lists client IP addresses involved in frequent malware connections, as identified by the L4 Traffic Monitor (L4TM). A computer that connects frequently to malware sites may be infected with malware that is trying to connect to a central command and control server and should be disinfected.

Figure 5-12 shows the Client Malware Risk page.

Figure 5-12 Client Malware Risk Page

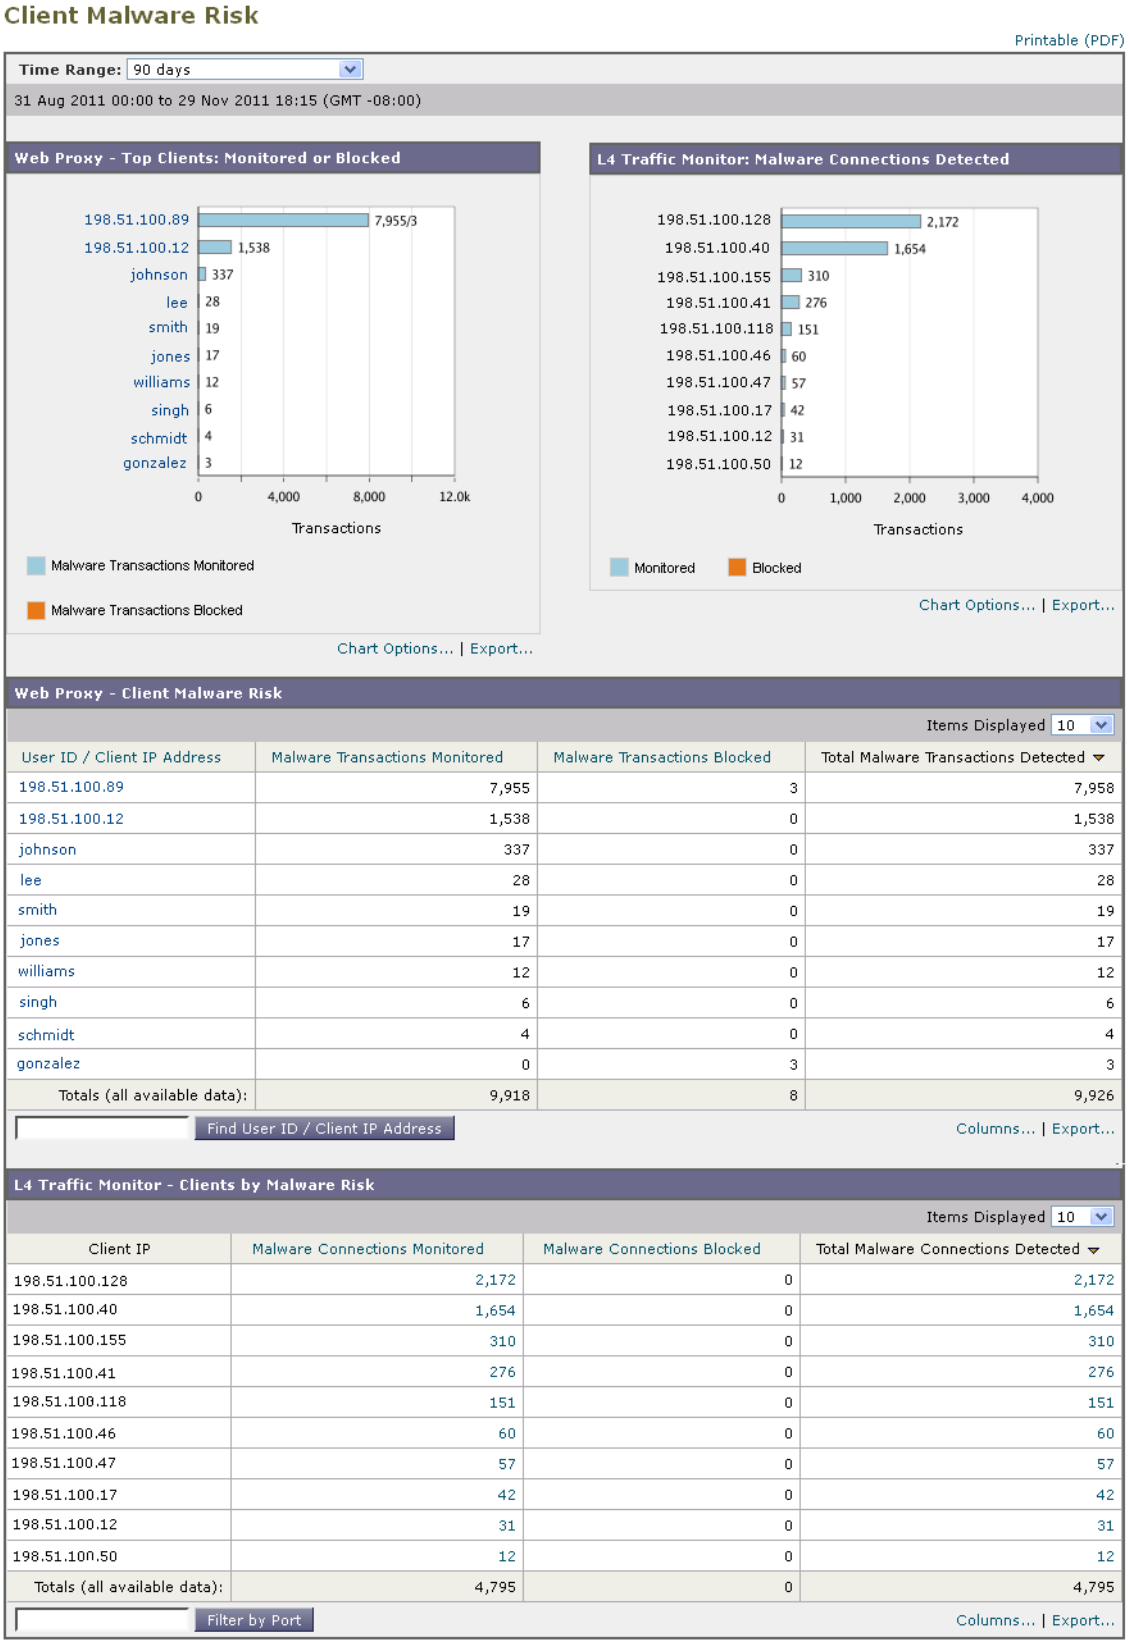

Table 5-11 describes the information on the Client Malware Risk page.

Table 5-11 Client Malware Risk Report Page Components

Section	Description
Time Range (drop-down list)	A menu that allows you to choose the time range of the data contained in the report. For more information, see Choosing a Time Range for Reports, page 3-4.
Web Proxy: Top Clients Monitored or Blocked	This chart displays the top ten users that have encountered a malware risk.
L4 Traffic Monitor: Malware Connections Detected	This chart displays the IP addresses of the ten computers in your organization that most frequently connect to malware sites. This chart is the same as the "Top Client IPs" chart on the L4 Traffic Monitor Report, page 5-42. See that section for more information and chart options.
Web Proxy: Client Malware Risk	The Web Proxy: Client Malware Risk table shows detailed information about particular clients that are displayed in the Web Proxy: Top Clients by Malware Risk section. You can click each user in this table to view the User Details page associated with that client. For information about that page, see the User Details (Web Reporting), page 5-19. Clicking on any of the links in the table allows you to view more granular details about individual users and what activity they are performing that is triggering the malware risk. For example, clicking on the link in the "User ID / Client IP Address" column takes you to a User page for that user.
L4 Traffic Monitor: Clients by Malware Risk	This table displays IP addresses of computers in your organization that frequently connect to malware sites. This table is the same as the "Client Source IPs" table on the L4 Traffic Monitor Report, page 5-42. For information about working with this table, see that section.

Tip To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

Web Reputation Filters Report

The **Web > Reporting > Web Reputation Filters** is a security-related reporting page that allows you to view the results of your set Web Reputation filters for transactions during a specified time range.

What are Web Reputation Filters?

Web Reputation Filters analyze web server behavior and assign a reputation score to a URL to determine the likelihood that it contains URL-based malware. It helps protect against URL-based malware that threatens end-user privacy and sensitive corporate information. The Web Security appliance uses URL reputation scores to identify suspicious activity and stop malware attacks before they occur. You can use Web Reputation Filters with both Access and Decryption Policies.

Web Reputation Filters use statistically significant data to assess the reliability of Internet domains and score the reputation of URLs. Data such as how long a specific domain has been registered, or where a web site is hosted, or whether a web server is using a dynamic IP address is used to judge the trustworthiness of a given URL.

The web reputation calculation associates a URL with network parameters to determine the probability that malware exists. The aggregate probability that malware exists is then mapped to a Web Reputation Score between -10 and +10, with +10 being the least likely to contain malware.

Example parameters include the following:

- URL categorization data
- Presence of downloadable code
- Presence of long, obfuscated End-User License Agreements (EULAs)
- Global volume and changes in volume
- Network owner information
- History of a URL
- Age of a URL
- Presence on any block lists
- Presence on any allow lists
- URL typos of popular domains
- Domain registrar information
- IP address information

For more information on Web Reputation Filtering, see ‘Web Reputation Filters’ in *Cisco IronPort AsyncOS for Web Security User Guide*.

Figure 5-13 Web Reputation Filters Page

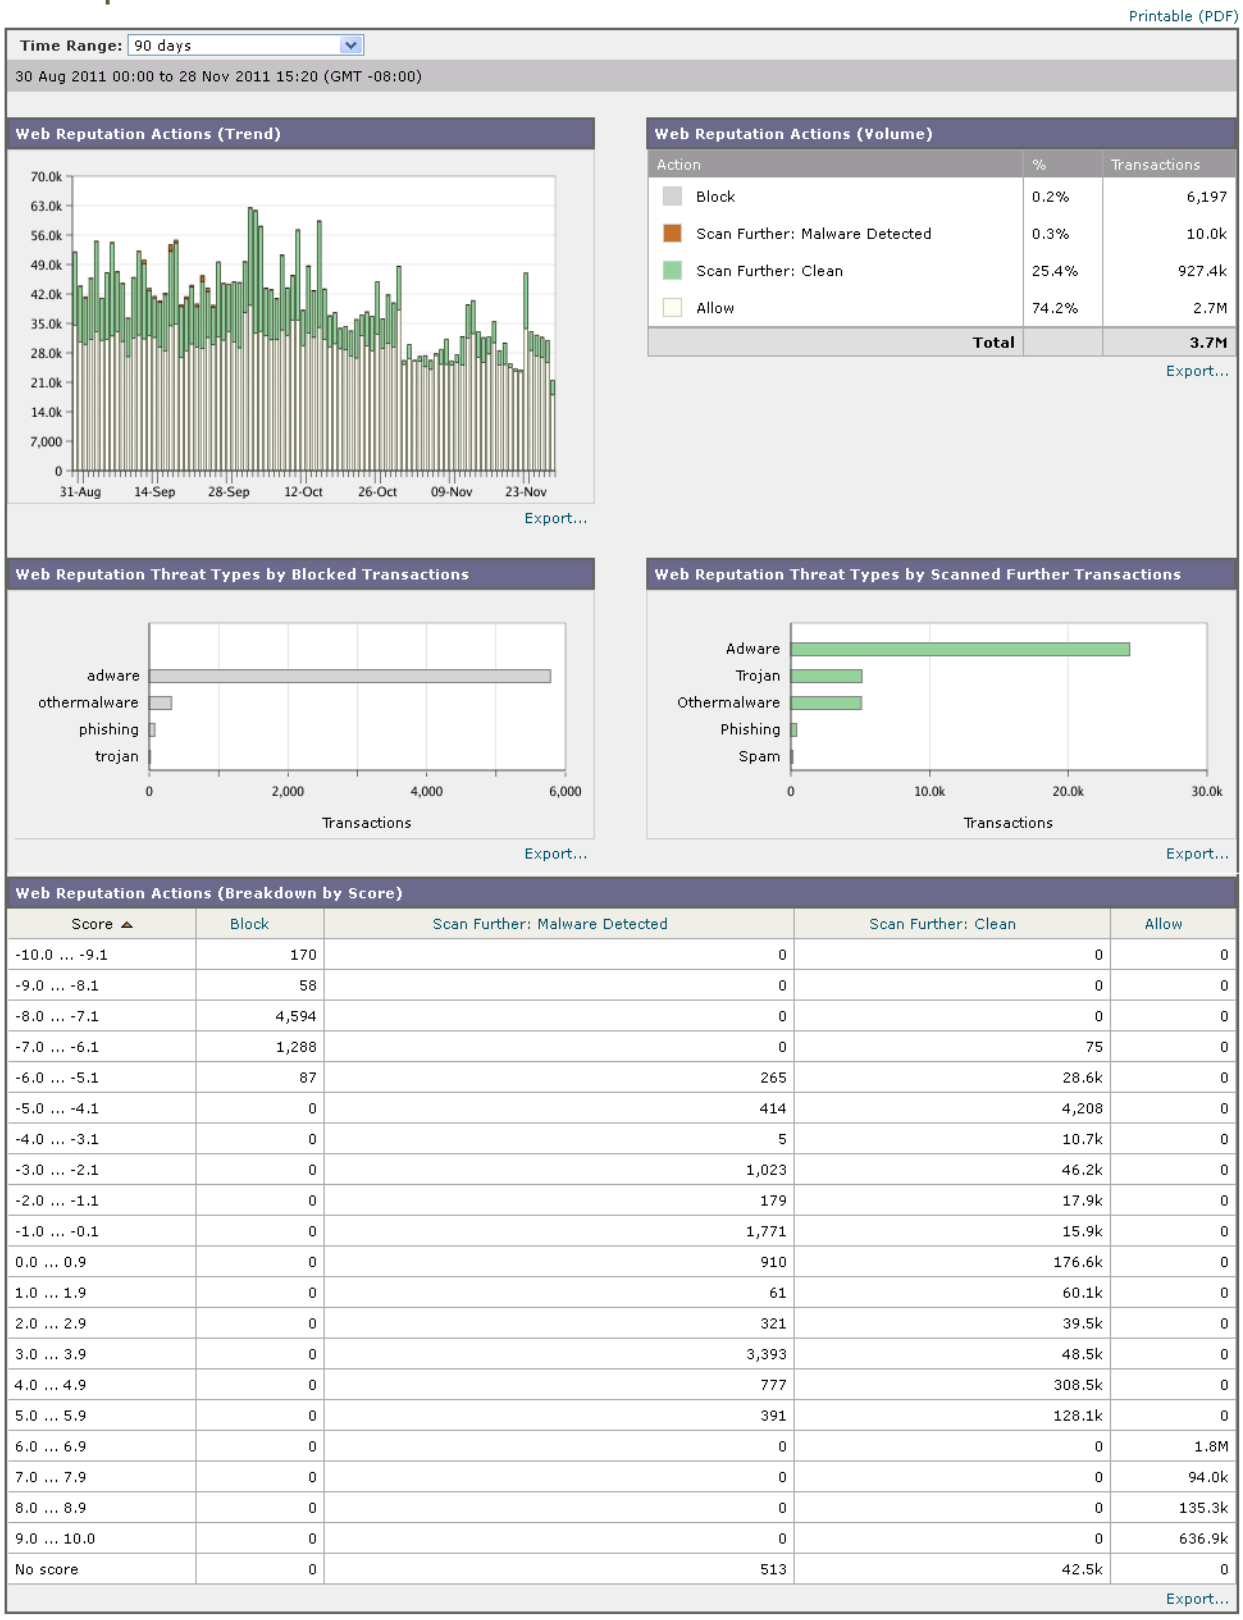

From the **Web Reputation Filters** page, you can view the following information:

Table 5-12 Details on the Web > Reporting > Web Reputation Filters 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" section on page 3-4.
Web Reputation Actions (Trend)	This section, in graph format, displays the total number of web reputation actions (vertical) against the time specified (horizontal timeline). From this you can see potential trends over time for web reputation actions.
Web Reputation Actions (Volume)	This section displays the web reputation action volume in percentages by transactions.
Web Reputation Threat Types by Blocked Transactions	This section displays the Web Reputation type that has been blocked.
Web Reputation Threat Types by Scanned Further Transactions	If Adaptive Scanning is enabled, this section displays the number of potentially threatening transactions caught. If Adaptive Scanning is not enabled, this section displays the Web Reputation type that has been blocked and due to this action, needs to be scanned further. If the result of Web Reputation filtering is to 'Scan Further', the transaction is passed to the Anti-Malware tool for additional scanning.
Web Reputation Actions (Breakdown by Score)	If Adaptive Scanning is not enabled, this interactive table displays the Web Reputation scores broken down for each action.

Tip To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

Adjusting Web Reputation Settings

Based on your report results, you may want to adjust the configured web reputation settings, for example adjust the threshold scores or enable or disable Adaptive Scanning. For specific information about configuring web reputation settings, see the user guide for your version of Cisco IronPort AsyncOS for Web Security.

L4 Traffic Monitor Report

The **Web > Reporting> L4 Traffic Monitor** page is a is a security-related reporting page that displays information about malware ports and malware sites that the L4 Traffic Monitors on your Web Security appliances have detected during the specified time range. It also displays IP addresses of clients that frequently encounter malware sites.

The L4 Traffic Monitor listens to network traffic that comes in over all ports on each Web Security appliance and matches domain names and IP addresses against entries in its own database tables to determine whether to allow incoming and outgoing traffic.

You can use data in this report to determine whether to block a port or a site, or to investigate why a particular client IP address is connecting unusually frequently to a malware site (for example, this could be because the computer associated with that IP address is infected with malware that is trying to connect to a central command and control server.)

Figure 5-14 shows the L4 Traffic Monitor page.

Tip To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

Figure 5-14 L4 Traffic Monitor Page

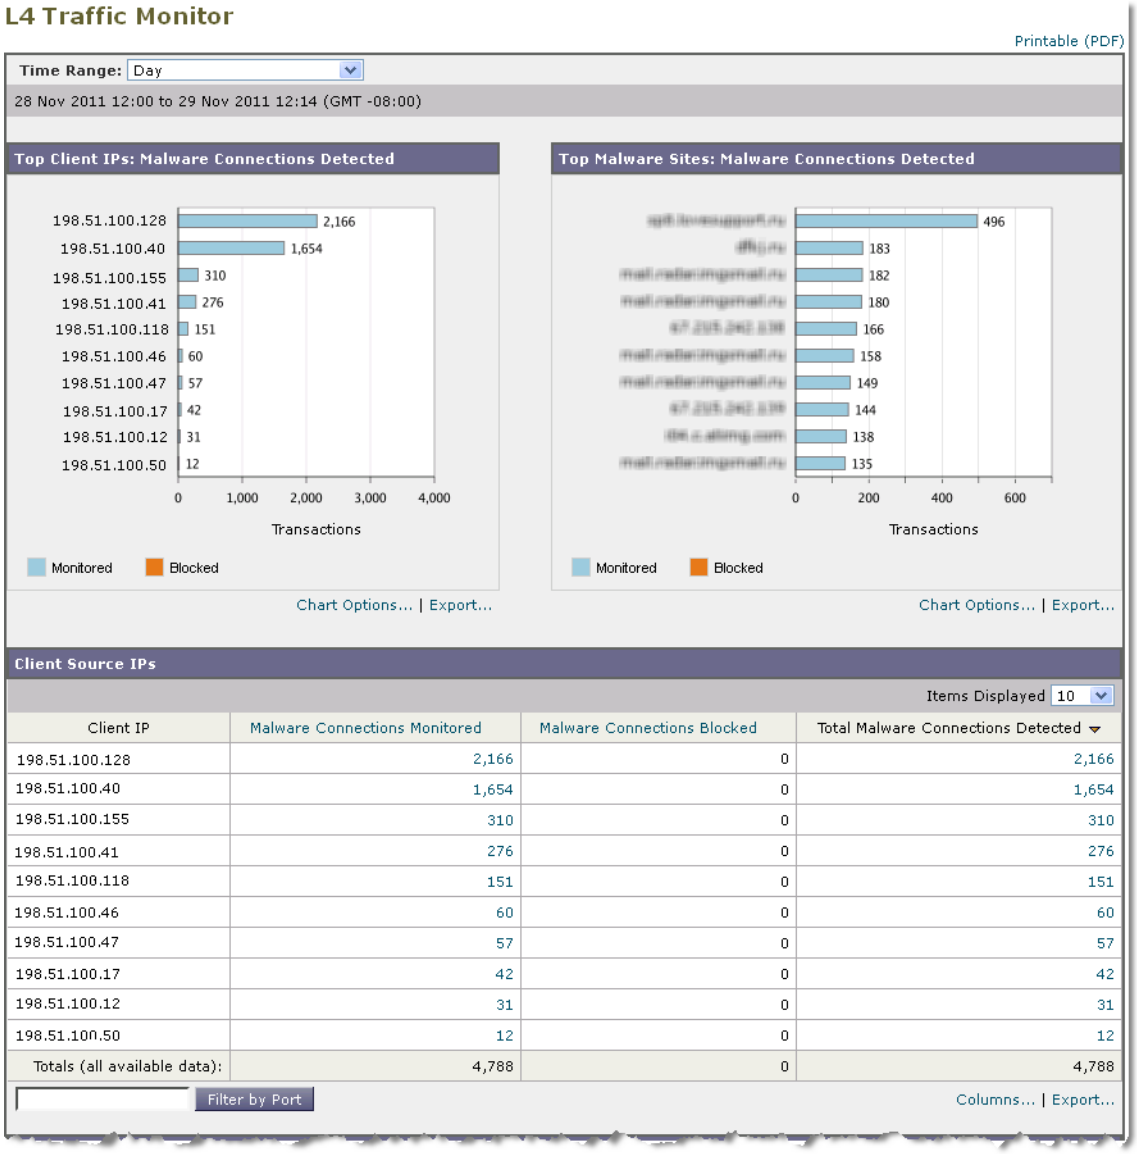

(Illustration continues on next page)

(Illustration continued from previous page)

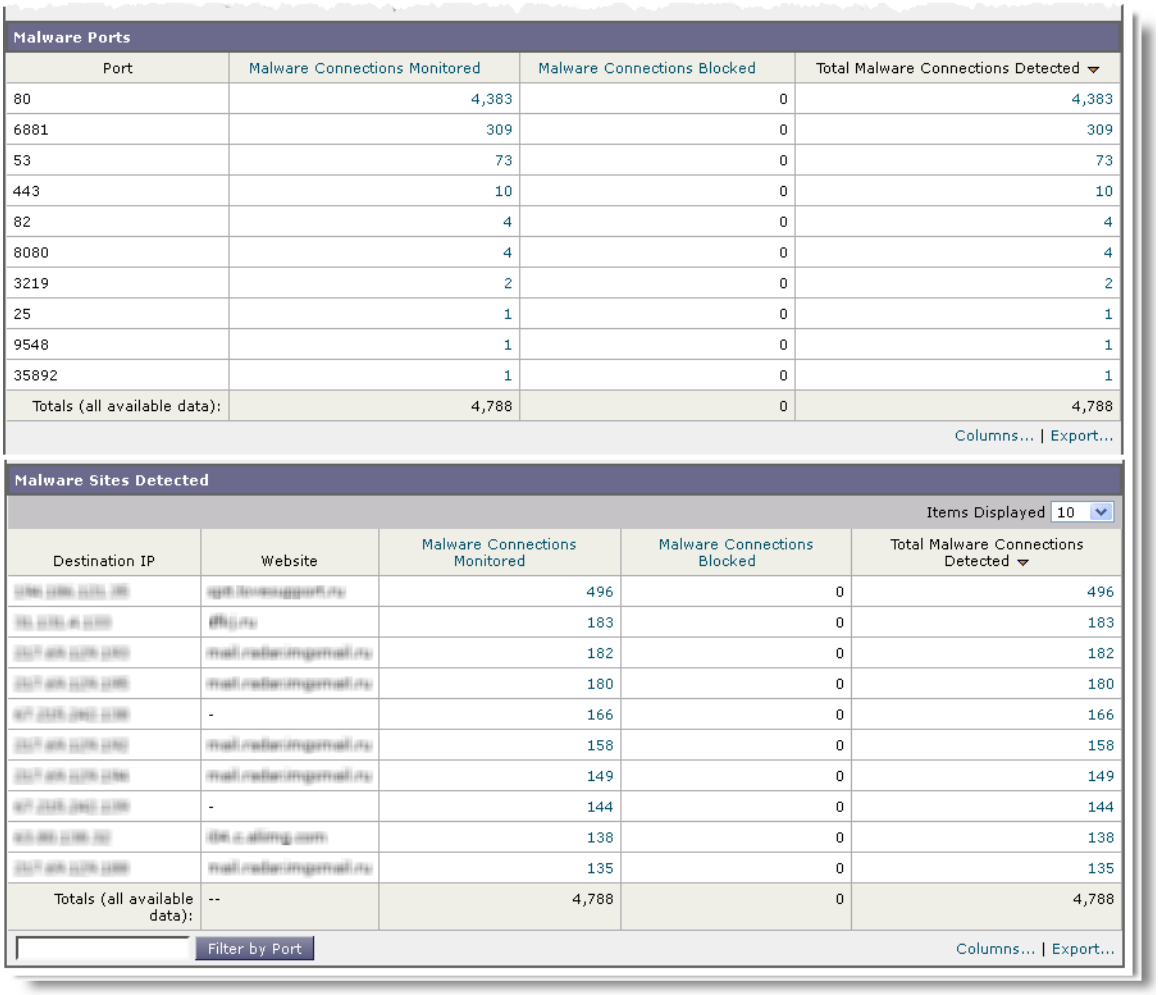

Malware Ports

Port	Malware Connections Monitored	Malware Connections Blocked	Total Malware Connections Detected
80	4,383	0	4,383
6881	309	0	309
53	73	0	73
443	10	0	10
82	4	0	4
8080	4	0	4
3219	2	0	2
25	1	0	1
9548	1	0	1
35892	1	0	1
Totals (all available data):	4,788	0	4,788

Columns... | Export...

Malware Sites Detected

Items Displayed 10

Destination IP	Website	Malware Connections Monitored	Malware Connections Blocked	Total Malware Connections Detected
[redacted]	[redacted]	496	0	496
[redacted]	[redacted]	183	0	183
[redacted]	[redacted]	182	0	182
[redacted]	[redacted]	180	0	180
[redacted]	-	166	0	166
[redacted]	[redacted]	158	0	158
[redacted]	[redacted]	149	0	149
[redacted]	-	144	0	144
[redacted]	[redacted]	138	0	138
[redacted]	[redacted]	135	0	135
Totals (all available data):	--	4,788	0	4,788

Filter by Port

Columns... | Export...

Table 5-13 describes the information on the L4 Traffic Monitor page.

Table 5-13 L4 Traffic Monitor Report Page Components

Section	Description
Time Range (drop-down list)	A menu that allows you to choose a time range on which to report. For more information, see Choosing a Time Range for Reports, page 3-4.
Top Client IPs	This section displays, in graph format, the IP addresses of computers in your organization that most frequently connect to malware sites. Click the Chart Options link below the chart to change the display from total Malware Connections Detected to Malware Connections Monitored or Malware Connections Blocked. This chart is the same as the "L4 Traffic Monitor: Malware Connections Detected" chart on the Client Malware Risk Report, page 5-37.

Table 5-13 L4 Traffic Monitor Report Page Components (continued)

Section	Description
Top Malware Sites	This section displays, in graph format, the top malware domains detected by the L4 Traffic Monitor. Click the Chart Options link below the chart to change the display from total Malware Connections Detected to Malware Connections Monitored or Malware Connections Blocked.
Client Source IPs	This table displays the IP addresses of computers in your organization that frequently connect to malware sites. To include only data for a particular port, enter a port number into the box at the bottom of the table and click Filter by Port. You can use this feature to help determine which ports are used by malware that "calls home" to malware sites. To view details such as the port and destination domain of each connection, click an entry in the table. For example, if one particular client IP address has a high number of Malware Connections Blocked, click the number in that column to view a list of each blocked connection. The list is displayed as search results in the L4 Traffic Monitor tab on the Web > Reporting > Web Tracking page. For more information about this list, see Searching for Transactions Processed by the L4 Traffic Monitor, page 5-54. This table is the same as the "L4 Traffic Monitor - Clients by Malware Risk" table on the Client Malware Risk Report, page 5-37.
Malware Ports	This table displays the ports on which the L4 Traffic Monitor has most frequently detected malware. To view details, click an entry in the table. For example, click the number of Total Malware Connections Detected to view details of each connection on that port. The list is displayed as search results in the L4 Traffic Monitor tab on the Web > Reporting > Web Tracking page. For more information about this list, see Searching for Transactions Processed by the L4 Traffic Monitor, page 5-54.
Malware Sites Detected	This table displays the domains on which the L4 Traffic Monitor most frequently detects malware. To include only data for a particular port, enter a port number into the box at the bottom of the table and click Filter by Port. You can use this feature to help determine whether to block a site or a port. To view details, click an entry in the table. For example, click the number of Malware Connections Blocked to view the list of each blocked connection for a particular site. The list is displayed as search results in the L4 Traffic Monitor tab on the Web > Reporting > Web Tracking page. For more information about this list, see Searching for Transactions Processed by the L4 Traffic Monitor, page 5-54.

Tip To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

Related Topics

- Searching for Transactions Processed by the L4 Traffic Monitor, page 5-54

SOCKS Proxy Report

The **Web > Reporting > SOCKS Proxy** Page allows you to view data and trends for transactions processed through the SOCKS proxy, including information about destinations and users.

Note The destination shown in the report is the address that the SOCKS client (typically a browser) sends to the SOCKS proxy.

To change SOCKS policy settings, see the *Cisco IronPort AsyncOS for Web Security User Guide*.

SOCKS Proxy

Printable (PDF)

Time Range: Day

31 Oct 2012 15:00 to 01 Nov 2012 15:56 (GMT -07:00)

Top Destinations for SOCKS: Total Transactions

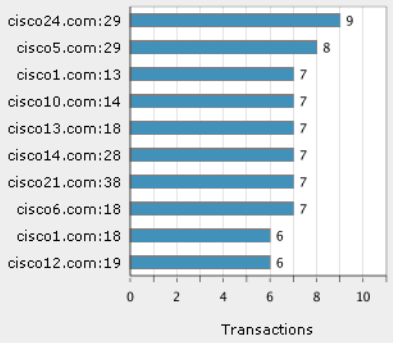

Chart Options... | Export...

Top Users for SOCKS: Total Transactions

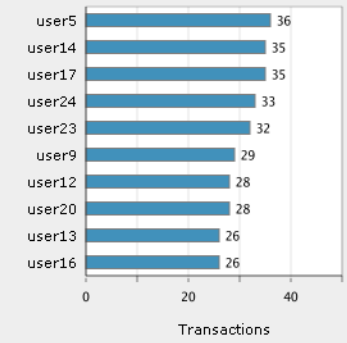

Chart Options... | Export...

Destinations

Items Displayed 10

Domain/IP:Port	TCP / UDP	Bandwidth Used	Transactions Allowed	Transactions Blocked	Total Transactions
cisco1.com:3	TCP	120B	4	0	4
cisco1.com:9	TCP	360B	4	0	4
cisco1.com:13	TCP	910B	7	0	7
cisco1.com:14	UDP	420B	3	0	3
cisco1.com:18	UDP	1,080B	6	0	6
cisco1.com:19	TCP	570B	3	0	3
cisco1.com:23	TCP	230B	1	0	1
cisco1.com:24	UDP	960B	4	0	4
cisco10.com:13	UDP	130B	1	0	1
cisco10.com:14	TCP	980B	7	0	7
Totals (all available data):	--	162.0KB	627	0	627

Find Domain/IP

Columns... | Export...

Users

Items Displayed 10

User ID or Client IP	Bandwidth Used	Transactions Allowed	Transactions Blocked	Total Transactions
user10	4,150B	17	0	17
user11	5,040B	20	0	20
user12	5,690B	28	0	28
user13	6,990B	26	0	26
user14	10,050B	35	0	35
user15	3,800B	13	0	13
user16	7,990B	26	0	26
user17	10,110B	35	0	35
user18	7,440B	25	0	25
user19	8,240B	24	0	24
Totals (all available data):	162.0KB	627	0	627

Find User ID or Client IP

Columns... | Export...

Related Topics

- Searching for Transactions Processed by the SOCKS Proxy, page 5-55

Reports by User Location

The **Web > Reporting > Reports by User Location** Page allows you to find out what activities that your mobile users are conducting from their local or remote systems.

Activities include:

- URL Categories that are being accessed by the local and remote users.
- Anti-Malware activity that is being triggered by sites the local and remote users are accessing.
- Web Reputation of the sites being accessed by the local and remote users.
- Applications that are being accessed by the local and remote users.
- Users (local and remote).
- Domains accessed by local and remote users.

Figure 5-15 Reports by User Location Page

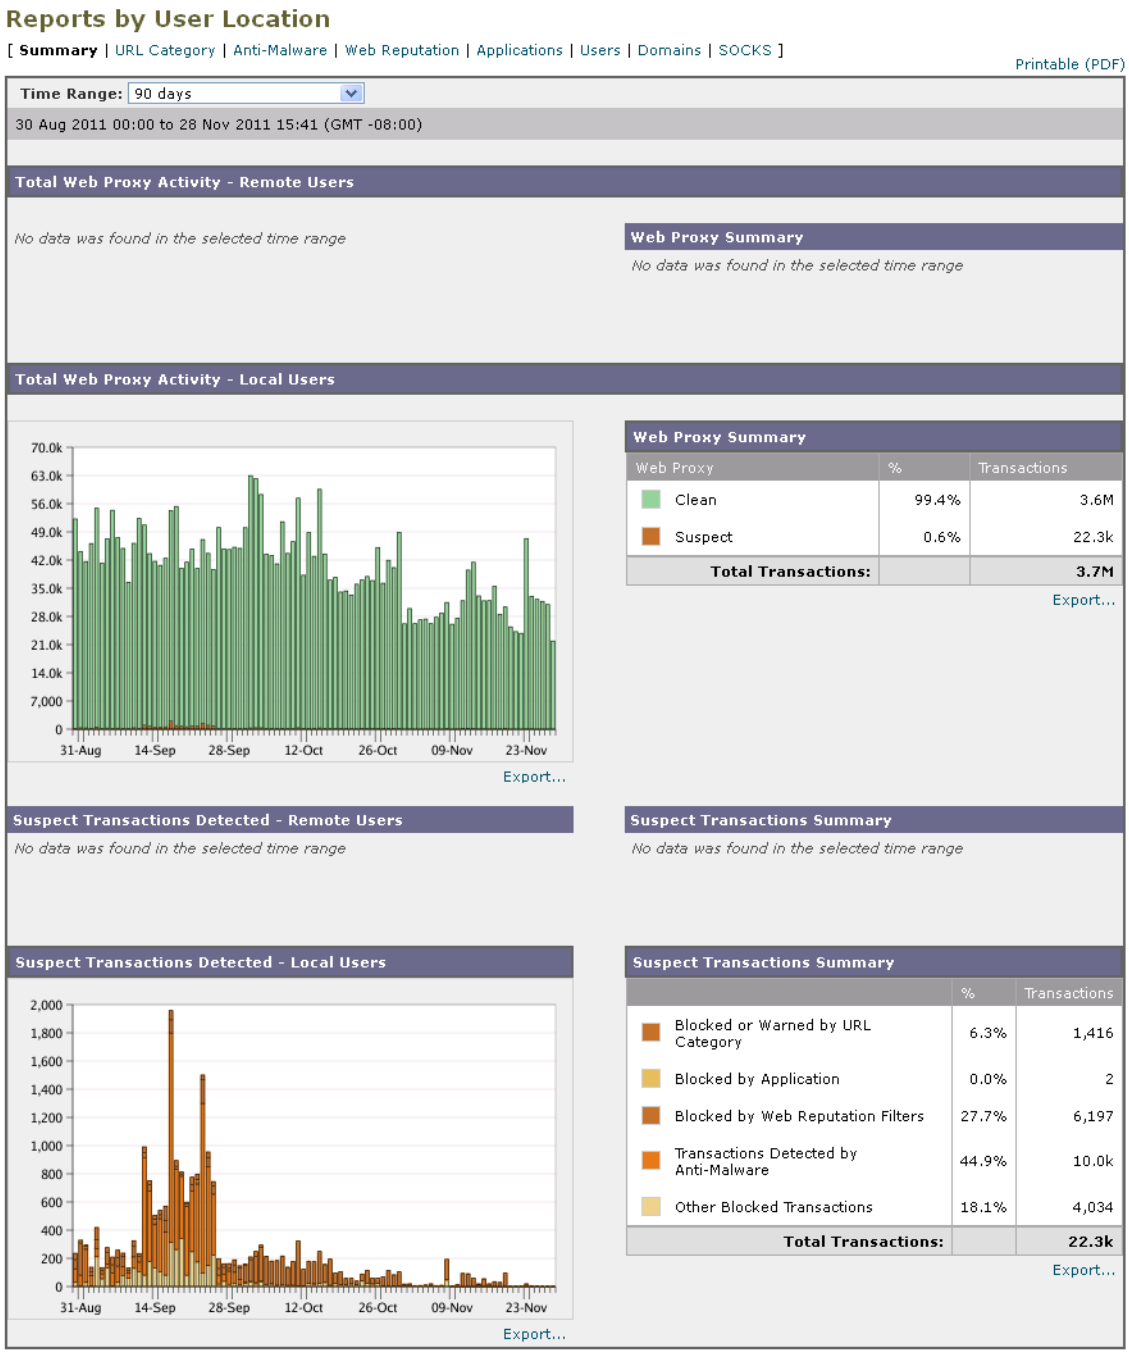

From the **Reports by User Location** page, you can view the following information:

Table 5-14 Details on the Web > Reporting > Reports by User Location 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" section on page 3-4.
Total Web Proxy Activity: Remote Users	This section displays, in graph format, the activity of your remote users (vertical) over the specified time (horizontal).
Web Proxy Summary	This section displays a summary of the activities of the local and remote users on your system.
Total Web Proxy Activity: Local Users	This section displays, in graph format, the activity of your remote users (vertical) over the specified time (horizontal).
Suspect Transactions Detected: Remote Users	This section displays, in graph format, the suspect transactions that have been detected due to access policies that you have defined for your remote users (vertical) over the specified time (horizontal).
Suspect Transactions Summary	This section displays a summary of suspected transactions of the remote users on your system.
Suspect Transactions Detected: Local Users	This section displays, in graph format, the suspect transactions that have been detected due to access policies that you have defined for your remote users (vertical) over the specified time (horizontal).
Suspect Transactions Summary	This section displays a summary of suspected transactions of the local users on your system.

From the **Reports by User Location** page you can generate reports showing the activity of local and remote users. This allows you to easily compare local and remote activities of your users.

Tip To customize your view of this report, see Working with Interactive Web Reporting Pages, page 5-7.

Note You can generate a scheduled report for information on the Reports by User Location page. For information on scheduling a report, see the "About Scheduled and On-Demand Web Reports" section on page 5-61.

Web Tracking

Use the Web Tracking page to search for and view details about individual transactions or patterns of transactions that may be of concern. Depending on the services that your deployment uses, search in relevant tabs:

- Searching for Transactions Processed by Web Proxy Services, page 5-51
- Searching for Transactions Processed by the L4 Traffic Monitor, page 5-54
- Searching for Transactions Processed by the SOCKS Proxy, page 5-55

For more information about the distinction between the Web Proxy and the L4 Traffic Monitor, see the "Understanding How the Web Security Appliance Works" section in the *Cisco IronPort AsyncOS for Web Security User Guide*.

Related Topics

- About Web Tracking and Upgrades, page 5-55

Searching for Transactions Processed by Web Proxy Services

Use the **Proxy Services** tab on the **Web > Reporting > Web Tracking** page to search web tracking data aggregated from individual security components and acceptable use enforcement components. This data does not include L4 Traffic Monitoring data.

You might want to use it to assist the following roles:

- **HR or Legal manager.** Run an investigative report for an employee during a specific time period.

  For example, you can use the Proxy Services tab to retrieve information about a specific URL that a user is accessing, what time the user visited that URL, whether that URL is allowed, etc.

- **Network security administrator.** Examine whether the company network is being exposed to malware threats through employees' smartphones.

You can view search results for the transactions recorded (including blocked, monitored, warned, and completed) during a particular time period. You can also filter the data results using several criteria, such as URL category, malware threat, and application.

Note The Web Proxy only reports on transactions that include an ACL decision tag other than "OTHER-NONE.

For an example of Web Tracking usage, see the "Example 1: Investigating a User" section on page D-1.

For an example of how the Proxy Services tab can be used with other web reporting pages, see the "Using The URL Categories Page in Conjunction with Other Reporting Pages" section on page 5-27.

To search for instances of web activity that concern you:

Procedure

Step 1 On the Security Management appliance window, choose **Web > Reporting > Web Tracking**.

Step 2 Click the **Proxy Services** tab.

Step 3 To see all search and filtering options, click **Advanced**.

Step 4 Enter search criteria:

Table 5-15 Web Tracking Search Criteria on the Proxy Services Tab

Option	Description
Default Search Criteria	
Time Range	Choose the time range on which to report. For information on time ranges available on the Security Management appliance, see the "Choosing a Time Range for Reports" section on page 3-4.

Table 5-15 Web Tracking Search Criteria on the Proxy Services Tab

Option	Description
User/Client IP	Optionally, enter an authentication username as it appears in reports or a client IP address that you want to track. You can also enter an IP range in CIDR format, such as 172.16.0.0/16. When you leave this field empty, the search returns results for all users.
Website	Optionally, enter a website that you want to track. When you leave this field empty, the search returns results for all websites.
Transaction Type	Choose the type of transactions that you want to track, either All Transactions, Completed, Blocked, Monitored, or Warned.
Advanced Search Criteria	
URL Category	To filter by a URL category, select **Filter by URL Category** and type the first letter of a custom or predefined URL category by which to filter. Choose the category from the list that appears . . If the set of URL categories has been updated, some categories may be labeled "Deprecated". These categories are no longer being used for new transactions on at least one managed Web Security appliance. However, you can still search for recent transactions that occurred while the category was active. For more information about URL category set updates, see URL Category Set Updates and Reports, page 5-26. All recent transactions that match the category name are included, regardless of the engine name noted in the drop-down list.
Application	To filter by an application, select **Filter by Application** and choose an application by which to filter. To filter by an application type, select **Filter by Application Type** and choose an application type by which to filter.
Policy	To filter by a policy group, select **Filter by Policy** and enter a policy group name by which to filter. Make sure that you have declared the policy on the Web Security appliance.
Malware Threat	To filter by a particular malware threat, select **Filter by Malware Threat** and enter a malware threat name by which to filter. To filter by a malware category, select **Filter by Malware Category** and choose a malware category by which to filter.

Table 5-15 Web Tracking Search Criteria on the Proxy Services Tab

Option	Description
WBRS	In the WBRS section, you can filter by Web-Based Reputation Score and by a particular web reputation threat. • To filter by web reputation score, select **Score Range** and select the upper and lower values by which to filter. Or, you can filter for websites that have no score by selecting **No Score**. • To filter by web reputation threat, select **Filter by Reputation Threat** and enter a web reputation threat by which to filter. For more information on WBRS scores, see the *Cisco IronPort AsyncOS for Web Security User Guide.*
AnyConnect Secure Mobility	To filter by remote or local access, select **Filter by User Location** and choose an access type. To include all access types, select **Disable Filter**. (In previous releases, this option was labeled Mobile User Security.)
User Request	To filter by transactions that were actually initiated by the user, select **Filter by User-Requested Transactions**. **Note:** When you enable this filter, the search results include "best guess" transactions.
Web Appliance	To filter by a specific Web appliance, click on the radio button next to **Filter by Web Appliance** and enter the Web appliance name in the text field. If you select **Disable Filter**, the search includes all Web Security appliances associated with the Security Management appliance.

Step 5 Click **Search**.

Understanding Web Tracking Search Results

By default, results are sorted by time stamp, with the most recent result at the top.

Figure 5-16 Web Tracking Search Results (Proxy Services Tab)

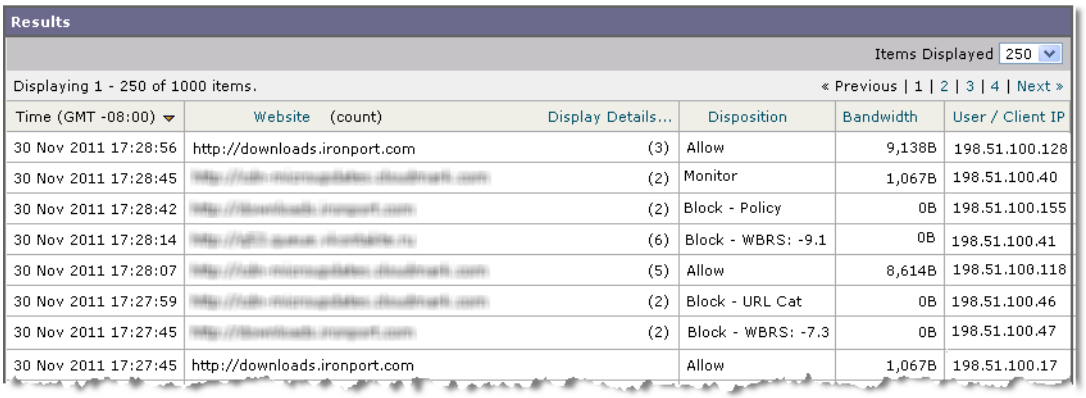

Results

Items Displayed 250

Displaying 1 - 250 of 1000 items. « Previous | 1 | 2 | 3 | 4 | Next »

Time (GMT -08:00)	Website (count)	Display Details...	Disposition	Bandwidth	User / Client IP
30 Nov 2011 17:28:56	http://downloads.ironport.com	(3)	Allow	9,138B	198.51.100.128
30 Nov 2011 17:28:45	[redacted]	(2)	Monitor	1,067B	198.51.100.40
30 Nov 2011 17:28:42	[redacted]	(2)	Block - Policy	0B	198.51.100.155
30 Nov 2011 17:28:14	[redacted]	(6)	Block - WBRS: -9.1	0B	198.51.100.41
30 Nov 2011 17:28:07	[redacted]	(5)	Allow	8,614B	198.51.100.118
30 Nov 2011 17:27:59	[redacted]	(2)	Block - URL Cat	0B	198.51.100.46
30 Nov 2011 17:27:45	[redacted]	(2)	Block - WBRS: -7.3	0B	198.51.100.47
30 Nov 2011 17:27:45	http://downloads.ironport.com		Allow	1,067B	198.51.100.17

In the Results window, you can view the following:

- The time the URL was accessed
- The website involved in the transaction
- The number of related transactions spawned by the user-initiated transaction, such as images loaded, javascripts run, and secondary sites accessed. This number appears in parentheses below the "Display Details" link in the column heading.
- The disposition (The result of the transaction. If applicable, shows the reason the transaction was blocked, monitored, or warned.)
- Bandwidth of the transaction
- User ID/Client IP address

Step 6 To view more information about the transactions, click the **Display Details...** link in the Website column heading.

Note If you need to view more than 1000 results, click the **Printable Download** link to obtain a CSV file that includes the complete set of raw data, excluding details of related transactions.

Tip If a URL in the results is truncated, you can determine the full URL by noting which host Web Security appliance processed the transaction, then checking the Accesslog on that appliance.

To view a list of up to 500 related transactions, click the **Related Transactions** link.

Searching for Transactions Processed by the L4 Traffic Monitor

The L4 Traffic Monitor tab on the **Web > Reporting > Web Tracking** page provides details about connections to malware sites and ports. You can search for connections to malware sites by the following types of information:

- Time range
- Site, using IP address or domain
- Port

- IP address associated with a computer in your organization
- Connection type
- The Web Security appliance that processed the connection

The first 1000 matching search results are displayed.

To view the hostname at the questionable site or the Web Security appliance that processed the transaction, click the Display Details link in the Destination IP Address column heading.

For more information about how you can use this information, see L4 Traffic Monitor Report, page 5-42.

Related Topics

- L4 Traffic Monitor Report, page 5-42

Searching for Transactions Processed by the SOCKS Proxy

You can search for transactions that meet a variety of criteria, including blocked or completed transactions; users; and destination domain, IP address, or port. You can also filter results by custom URL category, policy matched, and user location (local or remote).

Procedure

Step 1 Choose **Web > Reporting > Web Tracking**.

Step 2 Click the **SOCKS Proxy** tab.

Step 3 To filter results, click **Advanced**.

Step 4 Enter search criteria.

Step 5 Click **Search**.

Related Topics

- SOCKS Proxy Report, page 5-46

About Web Tracking and Upgrades

New web tracking 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 web tracking data and upgrades, see the Release Notes for your release.

System Capacity Page

The **Web > Reporting > System Capacity** page allows you to view the overall workload that is put on the Security Management appliance by the Web Security appliances. Most importantly, you can use the System Capacity page to track growth over time and plan for system capacity. Monitoring your Web Security appliances ensures that the capacity is appropriate to your volumes. Over time, volume inevitably rises and appropriate monitoring ensures that additional capacity or configuration changes can be applied proactively.

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

- Identify when Web Security appliances are exceeding recommended CPU capacity; this enables you to determine when configuration optimization or additional appliances are needed.
- For troubleshooting, identify which parts of the system are using the most resources.
- Identify response time and Proxy buffer memory.
- Identify the transactions per second, and any connections that are outstanding.

How to Interpret the Data You See on the 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.
- 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.

Note If you select **Year** for the time range for other reports, we recommend that you select the largest time range, 90 days.

To access the **System Capacity** page, perform the following:

Procedure

Step 1 On the Security Management appliance window, choose **Web > Reporting > System Capacity**.

Figure 5-17 System Capacity Page

System Capacity

Printable (PDF)

Time Range: 90 days

30 Aug 2011 00:00 to 28 Nov 2011 18:20 (GMT -08:00)

Overview of Averaged Usage and Performance

Web Security Appliance	CPU Usage %	Response Time (ms)	Proxy Buffer Memory (Bytes)	Transactions Per Second	Connections Out	Bandwidth Out (Bytes Per Second)
WSA_01	27.7%	511	0B	0	11	146
WSA_02	32.1%	523	0B	0	34	135
WSA_03	38.4%	541	0B	0	45	152

Columns... | Export...

Step 2 To view different types of data, click **Columns** and choose the data to view.

Step 3 To see the system capacity for a single appliance, click the appliance in the Web Security appliance column in the Overview of Averaged Usage and Performance table.

The System Capacity graphs appear for that appliance. The graphs on the page are divided into two sets:

- System Capacity - System Load
- System Capacity - Network Load

System Capacity - System Load

The first four graphs on the System Capacity window show the system load reports. These reports show the overall CPU usage on the appliances. AsyncOS is optimized to use idle CPU resources to improve transaction throughput. High CPU usage may not indicate a system capacity problem. If the high CPU usage is coupled with consistent, high-volume memory page swapping, you may have a capacity problem. This page also shows a graph that displays the amount of CPU used by different functions, including processing for the Web Security appliance reporting. The CPU-by-function graph is an 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.

Additionally, the Response Time/Latency and Transactions Per Second graphs shows the overall response time (in milliseconds), and transactions per second for the date range specified in the Time Range drop-down menu.

Figure 5-18 System Capacity - System Load

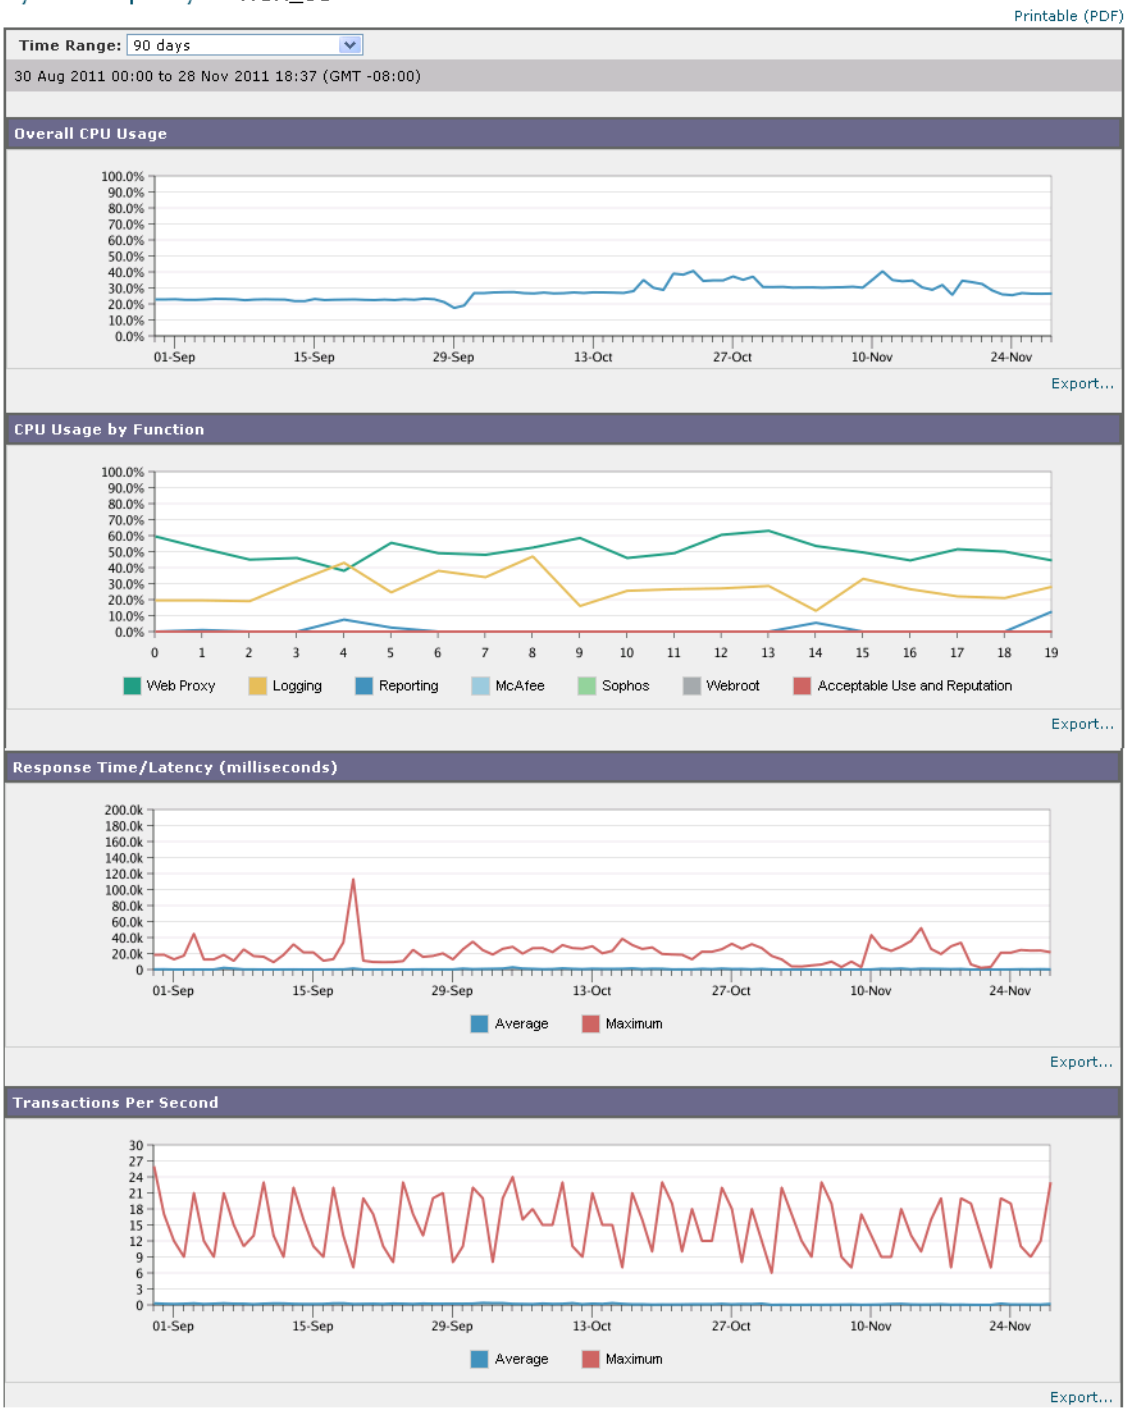

System Capacity - Network Load

The next graphs on the System Capacity window show the outgoing connections, the bandwidth out, and the proxy buffer memory statistics. You can view the results for a day, week, month, or year. It is important to understand the trends of normal volume and spikes in your environment.

The Proxy Buffer Memory may indicate spikes in network traffic during normal operation, but if the graph climbs steadily to the maximum, the appliance may be reaching its maximum capacity and you should consider adding capacity.

The following charts are on the same page as the charts in Figure 5-18, System Capacity - System Load, below those charts.

Figure 5-19 System Capacity - Network Load

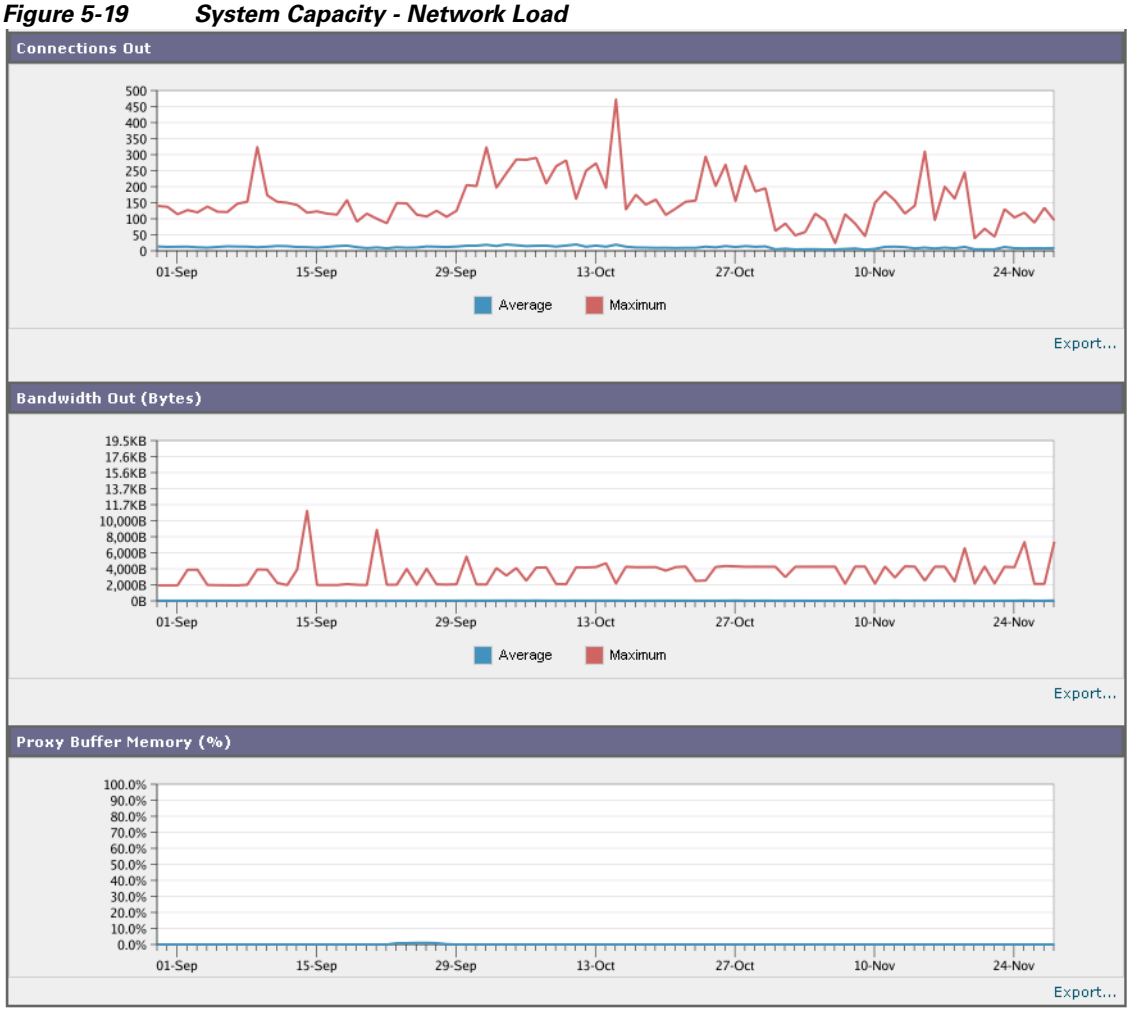

Note About Proxy Buffer Memory Swapping

The system is designed to swap proxy buffer memory regularly, so some proxy buffer memory swapping is expected and is not an indication of problems with your appliance. Unless the system *consistently* swaps proxy buffer memory in high volumes, proxy buffer memory swapping is normal and expected

behavior. If your system runs with extremely high volumes, and consistently swaps proxy buffer memory due to the high volumes, you may need to add Cisco IronPort appliances to your network or tune your configuration to ensure maximum throughput to improve performance.

Data Availability Page

The **Web > Reporting > Data Availability** page provides an overview of the date ranges for which reporting and web tracking data are available on the Security Management appliance for each managed Web Security appliance.

Figure 5-20 Web Reporting Data Availability Page

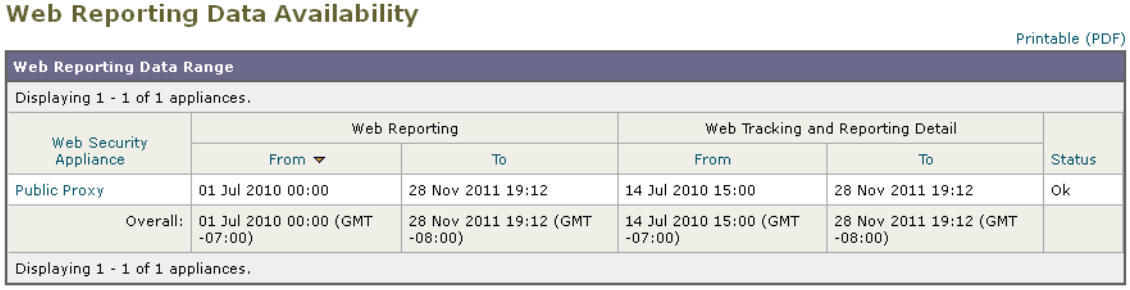

Note If Web Reporting is disabled, the Security Management appliance will not pull any new data from the Web Security appliance, but previously retrieved data is still present on the Security Management appliance. For information on how to manage disk usage, see the "Managing Disk Usage" section on page 13-50.

If the status is different between The Web Reporting 'From' and 'To' columns, and the Web Reporting and Tracking 'From' and 'To' columns, the most severe consequence appears in the Status column.

To view a graphical representation of data availability for a specific appliance, click an appliance in the Web Security appliance column.

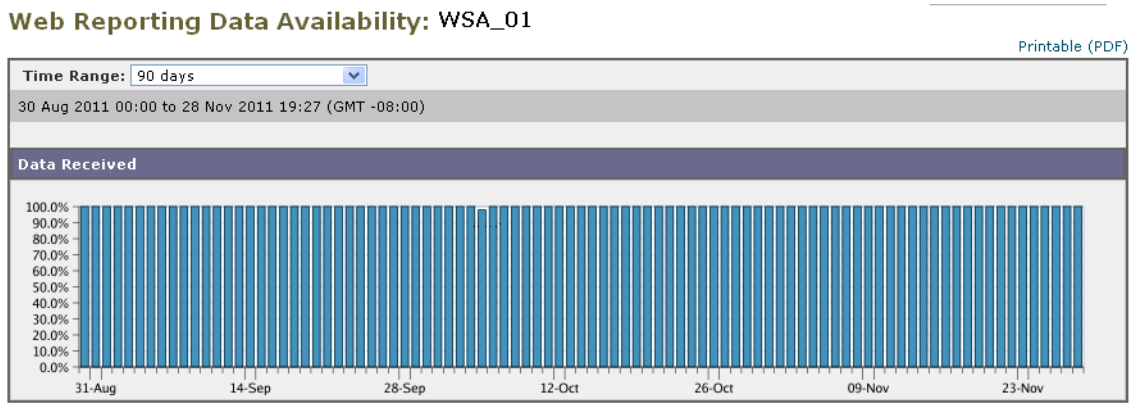

Note If Data Availability is used within a scheduled report for URL Categories, and there are gaps in data for any of the appliances, the following message is displayed at the bottom of the page: "Some data in this time range was unavailable."
If there are no gaps present, nothing appears.

About Scheduled and On-Demand Web Reports

Except as noted, you can generate the following types of Web Security reports either as scheduled or on-demand reports:

- Web Reporting Overview—For information on what is included on this page, see the "Web Reporting Overview" section on page 5-12.
- Users—For information on what is included on this page, see the "Users Report (Web)" section on page 5-16.
- Web Sites—For information on what is included on this page, see the "Web Sites Report" section on page 5-22.
- URL Categories—For information on what is included on this page, see the "URL Categories Report" section on page 5-24.
- Top URL Categories — Extended: For information on how to generate a report for Top URL Categories — Extended, see the Top URL Categories—Extended, page 5-63.

  This report is not available as an On-Demand report.
- Application Visibility—For information on what is included on this page, see the "Application Visibility Report" section on page 5-27.
- Top Application Types — Extended: For information on how to generate a report for Top URL Categories — Extended, see the Top Application Types—Extended, page 5-64.

  This report is not available as an On-Demand report.
- Anti-Malware—For information on what is included on this page, see the "Anti-Malware Report" section on page 5-31.
- Client Malware Risk—For information on what is included on this page, see the "Client Malware Risk Report" section on page 5-37.
- Web Reputation Filters—For information on what is included on this page, see the "Web Reputation Filters Report" section on page 5-39.
- L4 Traffic Monitor—For information on what is included on this page, see the "L4 Traffic Monitor Report" section on page 5-42.
- Mobile Secure Solution—For information on what is included on this page, see the "Reports by User Location" section on page 5-48.
- System Capacity—For information on what is included on this page, see the "System Capacity Page" section on page 5-55.

Scheduling Web Reports

This section includes the following:

- Adding Scheduled Reports, page 5-62
- Editing Scheduled Reports, page 5-63
- Deleting Scheduled Reports, page 5-63
- Additional Extended Reports, page 5-63

Note You can choose to make user names unrecognizable in all reports. For information, see Anonymizing User Names in Web Reports, page 5-4.

You can schedule reports to run on a daily, weekly, or monthly basis. Scheduled reports can be configured to include data for the previous day, previous seven days, previous month, previous calendar day (up to 250), 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).

The Security Management appliance retains the most recent reports that it generates — up to 1000 total versions for all reports. You can define as many recipients for reports as you want, including zero recipients. If you do not specify an email recipient, the system will still archive the reports. If you need to send the reports to a large number of addresses, however, you may want to create a mailing list instead of listing the recipients individually.

By default, the appliance archives the twelve most recent reports of each scheduled report. Reports are stored in the **/periodic_reports** directory of the appliance. (See Appendix A, "IP Interfaces and Accessing the Appliance" for more information.)

Adding Scheduled Reports

Procedure

Step 1 On the Security Management appliance window, choose **Web > Reporting > Scheduled Reports**.

Step 2 Click **Add Scheduled Report**.

Step 3 From drop-down menu next to **Type**, choose your report type.

Step 4 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 5 Choose the time range for the report from the **Time Range** drop-down menu.

Step 6 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 7 From the drop-down list next to **Number of Items**, choose the number of items that you want to be included in the generated report.

Valid values are from 2 through 20. The default value is 5.

Step 8 For **Charts**, click the default chart under **Data to display** and choose the data to display in each chart in the report.

Step 9 From the drop-down list next to **Sort Column**, select the column to sort the data by for this report. This allows you to create a scheduled report of Top 'N' items by any column available in the scheduled report.

Step 10 From the **Schedule** area, select the radio button next to the day, week, or month for your scheduled report.

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 address, the report is archived only.

Step 12 Click **Submit**.

Editing Scheduled Reports

To edit reports, go to the **Web > Reporting > Scheduled Reports** page and select the check boxes corresponding to the reports that you want to edit. Modify settings then click **Submit** to submit your changes on the page, then click the **Commit Changes** button to commit your changes on the appliance.

Deleting Scheduled Reports

To delete reports, go to the **Web > Reporting > Scheduled Reports** page and select the check boxes corresponding to the reports that you want to delete. To remove all scheduled reports, select the **All** check box, **Delete** and **Commit** your changes. Note that archived versions of deleted reports are not deleted.

Additional Extended Reports

Two additional reports are available only as Scheduled Reports on the Security Management appliance:

- Top URL Categories—Extended
- Top Application Types—Extended

Top URL Categories—Extended

The Top URL Categories —Extended report is useful for administrators who want to receive more detailed information than the URL Categories report can provide.

For example, in a typical URL Categories report, you can gather information measuring bandwidth usage by a particular employee at a larger URL Category level. To generate a more detailed report that monitors bandwidth usage for the top ten URLs for each URL Category, or top five users for each URL Category, use the Top URL Categories —Extended report.

Note

- The maximum number of reports that can be generated using this type of report is 20.
- Predefined URL category lists are occasionally updated. For more information about the impact of these updates on report results, see URL Category Set Updates and Reports, page 5-26.

To generate a Top URL Categories—Extended report, perform the following:

Procedure

Step 1 On the Security Management appliance window, choose **Web > Reporting > Scheduled Reports**.

Step 2 Click **Add Scheduled Report**.

Step 3 From the drop-down menu next to Type, choose **Top URL categories — Extended**.

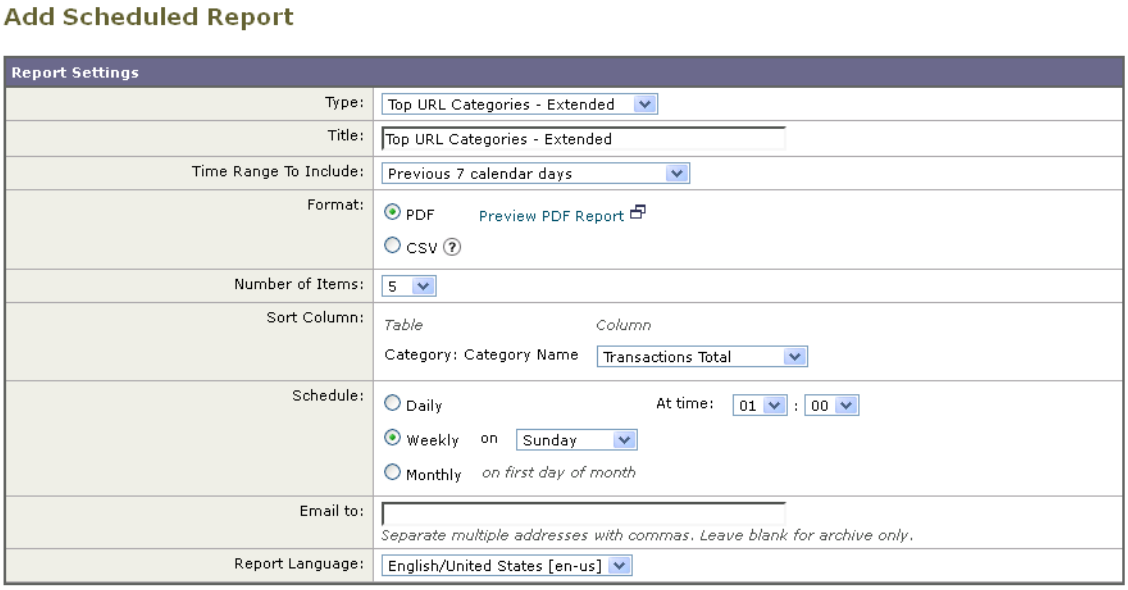

Step 4 In the **Title** text field, type the title of your URL extended report.

Step 5 Choose the time range for the report from the **Time Range** drop-down menu.

Step 6 Choose the format for the generated report.

The default format is PDF.

Step 7 From the drop-down list next to **Number of Items**, select the number of URL Categories that you want to be included in the generated report.

Valid values are from 2 through 20. The default value is 5.

Step 8 From the drop-down list next to Sort Column, select the column to sort the data by for this report. This allows you to create a scheduled report of Top 'N' items by any column available in the scheduled report.

Step 9 For **Charts**, click the default chart under **Data to display** and choose the data to display in each chart in the report.

Step 10 From the **Schedule** area, select the radio button next to the day, week, or month for your scheduled report.

Step 11 In the **Email** text field, type in the email address where the generated report will be sent.

Step 12 Click **Submit**.

Top Application Types—Extended

To generate a Top Application Type—Extended report, perform the following:

Procedure

Step 1 On the Security Management appliance window, choose **Web > Reporting > Scheduled Reports**.

Step 2 Click **Add Scheduled Report**.

Step 3 From the drop-down menu next to Type, choose **Top Application Types — Extended**.

The default options on the page will change.

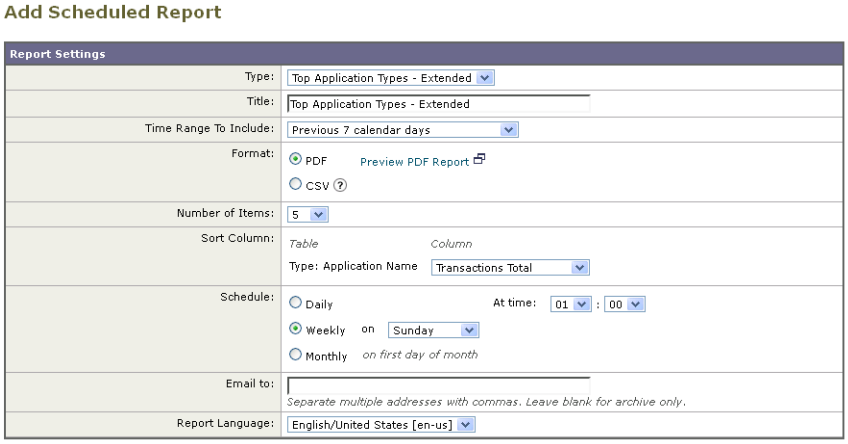

Step 4 In the **Title** text field, type the title of your report.

Step 5 Choose the time range for the report from the **Time Range** drop-down menu.

Step 6 Choose the format for the generated report.

The default format is PDF.

Step 7 From the drop-down list next to **Number of Items**, select the number of Application Types that you want to be included in the generated report.

Valid values are from 2 through 20. The default value is **5**.

Step 8 From the drop-down list next to **Sort Column**, select the type of column that you want to appear in the table. Choices include: Transactions Completed, Transactions Blocked, Transaction Totals.

Step 9 For **Charts**, click a default chart under **Data to display** and choose the data to display in each chart in the report.

Step 10 From the **Schedule** area, select the radio button next to the day, week, or month for your scheduled report.

Step 11 In the **Email** text field, type in the email address where the generated report will be sent.

Step 12 Click **Submit**.

Generating Web Reports on Demand

Most reports that you can schedule, you can also generate on demand.

Note Some reports are available only as Scheduled Reports, not on demand. See Additional Extended Reports, page 5-63.

To generate a report on demand, perform the following:

Procedure

Step 1 On the Security Management appliance window, choose, **Web > Reporting > Archived Reports**.

Step 2 Click on **Generate Report Now**.

Step 3 From the **Report type** section, choose a report type from the drop-down list.

The options on the page may change.

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

Step 6 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 7 Depending on the options available for the report, choose:

- **Number of rows**: The number of rows of data to display in the table.
- **Charts**: Which data to display in the chart(s) in the report:
- Click the default option under Data to display.
- **Sort Column**: The column to sort by for each table.

Step 8 From the Delivery Option section, choose the following:

- If you want this report to appear on the Archived Reports page, select the **Archive Report** checkbox.

Note Domain-Based Executive Summary reports cannot be archived.

- Check the **Email now to recipients** checkbox to email the report.
- In the text field, type in the recipient email addresses for the report.

Step 9 Click **Deliver This Report** to generate the report.

Viewing and Managing Archived Web Reports

The **Web > Reporting> Archived Reports** page lists the following:

- Reports that you schedule using the procedure in Adding Scheduled Reports, page 5-62
- Reports that you generate using the procedure in Generating Web Reports on Demand, page 5-65.

To view a report, click the report names in the Report Title column. The Show drop-down menu filters the types of reports that are listed on the **Archived Reports** page.

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.

The appliance stores up to 12 instances of each scheduled report (up to 1000 reports). Archived reports are stored in the `/periodic_reports` directory on the appliance. Archived reports are deleted automatically. As new reports are added, older reports are removed to keep the number at 1000. The limit of 12 instances applies to each scheduled report with the same name and time range.

CHAPTER 6

Tracking Email Messages

- Tracking Service Overview, page 6-1
- Setting Up Centralized Message Tracking, page 6-2
- Searching for Email Messages, page 6-4
- Understanding Tracking Query Results, page 6-7

Tracking Service Overview

The tracking service of the 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:

- **Envelope information:** Find messages from particular envelope senders or recipients by entering the text strings to match.
- **Subject header:** Match a text string in the subject line. Warning: Do not use this type of search in environments where regulations prohibit such tracking.
- **Time frame:** Find a message that was sent between specified dates and times.
- **Sender IP address or rejected connections:** Search for messages from a particular IP address, or show rejected connections in the search results.
- **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.

- **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.
- **Message ID:** Find messages by identifying the SMTP "Message-ID:" header or the Cisco IronPort message ID (MID).
- **Email Security appliance (host):** Narrow search criteria to particular Email Security appliances, or search across all managed appliances.

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, page 6-2
- Configuring Centralized Message Tracking on Email Security Appliances, page 6-2
- Adding the Centralized Message Tracking Service to Each Managed Email Security Appliance, page 6-3

Enabling Centralized Email Tracking on a Security Management Appliance

Procedure

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

Step 2 In the Message Tracking Service section, click **Enable**.

Step 3 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 4 **Submit** and commit your changes.

Configuring Centralized Message Tracking on Email Security Appliances

Procedure

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 *Cisco IronPort AsyncOS for Email Security Advanced Configuration Guide*.

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.

Procedure

Step 1 On the Security Management appliance, 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 Message Tracking** service.

Step 3 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 Cisco IronPort 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 4 Click **Submit**.

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

Step 6 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 DLP Information in Message Tracking, page 12-21.

Searching for Email Messages

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.

Procedure

Step 1 On the Security Management appliance window, 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.

  Note International character sets are not supported in the subject header.

- **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 Event:** Select the events to track. Options are Virus Positive, Spam Positive, Suspect Spam, Delivered, DLP Violations (you can enter the name of a DLP policy and select violation severities), Hard Bounced, Soft Bounced, Currently in Outbreak Quarantine, 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.

- **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.

Figure 6-1 Message Tracking Query Results

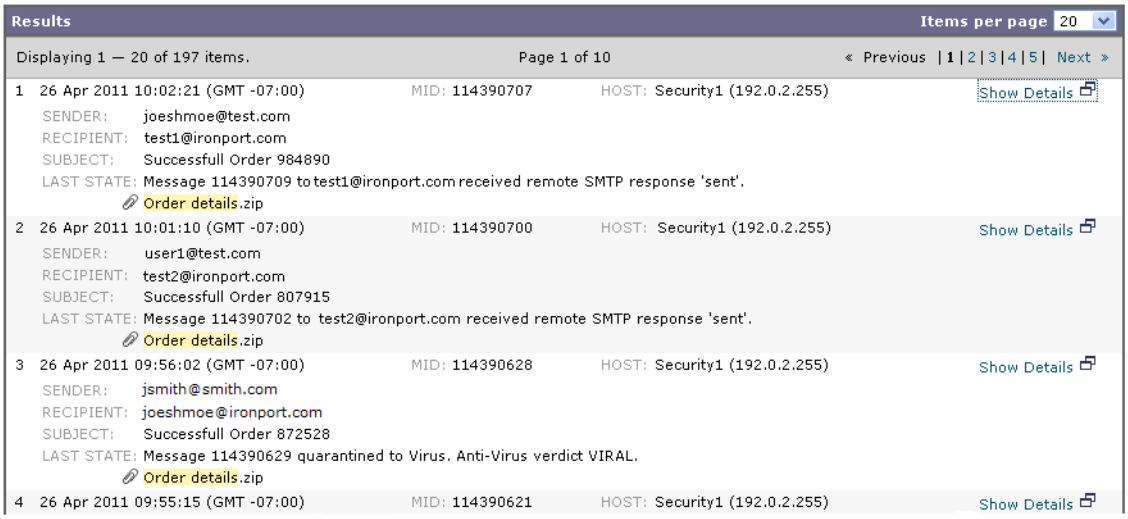

Your search criteria are highlighted in each row.

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.

Procedure

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 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 3 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.

Understanding Tracking Query Results

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 **Show Details** link for that message. For more information, see the "Message Details" section on page 6-7.

Note

- Messages with 50 or more recipients will not appear in tracking query results. This issue will be resolved in a future release.
- You can choose to display up to 1000 search results when you specify your query.
- 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 and Tracking Data, page 3-8.

Message Details

To view detailed information about a particular email message, including the message header information and processing details, click **Show Details** for any item in the search results list. A new window opens with the message details.

The message details include the following sections:

- Envelope and Header Summary, page 6-8
- Sending Host Summary, page 6-8

- Processing Details, page 6-8

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 IronPort 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.

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.

DLP Matched Content Tab

This section 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 DLP Information in Message Tracking, page 12-21.

CHAPTER 7

Managing the Cisco IronPort Spam Quarantine

- Understanding the Cisco IronPort Spam Quarantine, page 7-1
- Setting Up the Centralized Spam Quarantine, page 7-2
- Configuring Administrative User Access to the Cisco IronPort Spam Quarantine, page 7-7
- Configuring Spam Management Features for End Users, page 7-8
- Using End User Safelists and Blocklists, page 7-14
- Managing Messages in the Cisco IronPort Spam Quarantine, page 7-16

Understanding the Cisco IronPort Spam Quarantine

A Cisco IronPort Spam Quarantine holds spam and suspected spam messages for email users in your organization. The spam quarantine provides a safeguard mechanism for organizations that are concerned about "false positives" — that is, legitimate email messages that are quarantined or deleted as spam. This feature allows end users and administrators to review messages that are flagged as spam before making a final determination. In addition, if you enable the safelist/blocklist feature, end users (email users) can exercise control over which messages are marked as spam.

Note A *system quarantine* is distinct from a spam quarantine; it resides on an Email Security appliance and holds quarantined messages based on various actions that AsyncOS performs, such as content filtering, anti-virus scanning, and applying Outbreak Filters.

A *local* Cisco IronPort Spam Quarantine resides on an Email Security gateway appliance. You can also have messages sent to an *external* Cisco IronPort Spam Quarantine, which resides on a separate Cisco IronPort appliance — typically a Security Management appliance.

Note You can implement end user access to the Cisco IronPort Spam Quarantine only for specified users or groups of users. Also, after you initially implement end user access, you might later decide to disable access if end users rarely view and release messages in the quarantine.

You can configure AsyncOS to send a notification email to end users, informing them of quarantined spam and suspected spam messages. The notification contains a summary of the messages currently in the Cisco IronPort Spam Quarantine for that user. The user can view the messages and decide whether to have them delivered to the Email Inbox or delete them. Users can also search through their quarantined messages. Users can access the quarantine through the notification message, or they can access the

quarantine directly by using a web browser. (Direct end user access to the quarantine requires authentication. For more information, see the "Configuring End User Quarantine Access" section on page 7-8).

By default, the Cisco IronPort Spam Quarantine is self-maintaining. AsyncOS periodically deletes mail from the Cisco IronPort Spam Quarantine to prevent old messages from consuming all of the quarantine space.

All administrator-level users (such as the default admin user) can access and modify the Cisco IronPort Spam Quarantine. AsyncOS operator users and users to whom you assign Spam Quarantine access via a custom role can view and manage the quarantine content, but they cannot change the quarantine settings. Mail end users can access their own messages in the quarantine if you enable end user access to the Cisco IronPort Spam Quarantine.

Setting Up the Centralized Spam Quarantine

Before you centralize the spam quarantine, the local spam quarantine must be configured, working, and tested on your Email Security appliance.

To move the active spam quarantine from the Email Security appliance to the Security Management appliance, complete the following procedures in order:

- Identifying Required IP Addresses, page 7-2
- Configuring the Cisco IronPort Spam Quarantine Service on the Security Management Appliance, page 7-2
- Configuring Interfaces on the Security Management Appliance, page 7-4
- Configuring Email Security Appliances for Centralized Spam Quarantine, page 7-5
- Adding the Centralized Spam Quarantine Service to Each Managed Email Security Appliance, page 7-6

Identifying Required IP Addresses

Obtain or identify an IP address to use in the procedure in Configuring an Outbound IP Interface on the Security Management Appliance, page 7-4. This will typically be the Data 2 interface on the Security Management appliance. For more information about network requirements, see Appendix B, "Assigning Network and IP Addresses."

Configuring the Cisco IronPort Spam Quarantine Service on the Security Management Appliance

Procedure

Step 1 On the Security Management appliance window, choose **Management Appliance > Centralized Services > Spam Quarantine**

Step 2 If you are enabling the Cisco IronPort 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 3 If you are editing an existing configuration, click **Edit Settings** in the Cisco IronPort Spam Quarantine Settings section.

Step 4 In the **Quarantine IP Interface** section, specify the appropriate IP interface and port for the quarantine from the drop-down list.

By default, the 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.

Step 5 In the **Deliver Messages Via** section, type the primary and alternate routing for delivering mail in the corresponding text fields.

Because the Security Management appliance does not send messages directly, all outgoing quarantine-related email (such as spam notifications and messages released from 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 Data2 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.

Step 6 In the **Schedule Delete After** section, specify the number of days to hold messages before deleting them.

Alternatively, select the **Do not schedule a delete** radio button to disable scheduled deletions. It is recommended that you configure the quarantine to schedule deletions. When the quarantine fills to capacity, it deletes the oldest messages first.

Step 7 In the **Default Language** section, specify a default language.

This is the language that end users see when they access the Cisco IronPort Spam Quarantine.

Step 8 (Optional) In the **Notify Cisco IronPort upon Message Release**, check the checkbox to send a copy of released messages to Cisco IronPort for analysis.

It is recommended that you configure the quarantine to submit released messages for analysis.

Step 9 (Optional) In the **Spam Quarantine Appearance** section, customize the page that end users access when they view the quarantine.

Choices include:

- Use Current logo
- Use Cisco IronPort Spam Quarantine logo
- Upload Custom logo

  If you choose ‘Upload Custom logo’, the logo appears at the top of the Cisco IronPort Spam Quarantine page when the user logs in to view quarantined messages. The logo should be a .jpg, .gif, or .png file that is at most 550 x 50 pixels. If a logo file is not supplied, the default Cisco IronPort Spam Quarantine logo is used.

Step 10 (Optional) In the **Login Page Message** text field, type in a login page message. The message appears to end users when they are prompted to log in to the quarantine.

Step 11 Optionally, modify the list of users authorized to view the Cisco IronPort Spam Quarantine. For more information, see Configuring Administrative User Access to the Cisco IronPort Spam Quarantine, page 7-7.

Step 12 Optionally, configure end user access and spam notification settings. For more information, see Configuring Spam Management Features for End Users, page 7-8.

Step 13 Submit and commit your changes.

Configuring Interfaces on the Security Management Appliance

- Configuring an Outbound IP Interface on the Security Management Appliance, page 7-4
- Configuring the IP Interface for Spam Quarantine Access, page 7-4

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.

Procedure

Step 1 Use this procedure in conjunction with the information in Configuring IP Interfaces, page A-2.

Step 2 On the Security Management appliance, choose **Management Appliance > Network > IP Interfaces**.

Step 3 Click **Add IP Interface**.

Step 4 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 5 Submit and commit your changes.

Configuring the IP Interface for Spam Quarantine Access

When administrators and end users access the spam quarantine, a dedicated browser window opens.

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > Network > IP Interfaces**.

Step 2 Click the name of the **Management** interface.

Step 3 In the Spam Quarantine section, configure settings for access to the spam quarantine:

- Select HTTP and/or HTTPS and specify port(s).
- Specify the URL that appears in notifications and in the spam quarantine browser window.

  For example, if you do not want to expose the hostname of your Security Management appliance to end-users, you can specify an alternate hostname.

  Be sure your DNS server can resolve any URL or IP address that you enter here.

Step 4 Submit and commit your changes.

Configuring Email Security Appliances for Centralized Spam Quarantine

In order to use centralized Cisco IronPort Spam Quarantine services on the Security Management appliance, you must change some spam quarantine settings on the Email Security appliances:

Do This	See
Verify that spam quarantine is working properly on the Email Security appliance. If there are any issues, fix them before you centralize the spam quarantine.	—
Enable and configure the Email Security appliance to use the Security Management appliance as an external spam quarantine	Configuring the Email Security Appliance for Centralized Spam Quarantine, page 7-5
Disable the local spam quarantine on the Email Security appliance	"Disabling the Local Spam Quarantine" in the *Cisco IronPort AsyncOS for Email Security Daily Management Guide.* Ignore any warnings to adjust Mail Policies as a result of this change. Mail Policies will use the external spam quarantine.
Evaluate approaches for managing existing local spam quarantine messages on the Email Security appliance	"Migrating from a Local Spam Quarantine to an External Quarantine" in the *Cisco IronPort AsyncOS for Email Security Daily Management Guide.*

Configuring the Email Security Appliance for Centralized Spam Quarantine

If you want Email Security appliances to use the Cisco IronPort Spam Quarantine on the Security Management appliance, you need to configure the external quarantine settings on the Email Security appliances.

Note If a different external spam quarantine was previously configured for the Email Security appliances, first disable the external spam quarantine setting.

To configure the external quarantine settings, complete the following steps **on each** Email Security appliance:

Procedure

Step 1 On the Email Security appliance, navigate to **Security Services > External Spam Quarantine**.

Step 2 Click the **Configure** button.

Step 3 Select the check box to enable the external spam quarantine.

Step 4 Enter the name of the Cisco IronPort Spam Quarantine. Alternatively, you can enter the name of the Security Management appliance where the quarantine resides.

Step 5 Enter the IP address for the correct interface of the Security Management appliance:

Typically, this is the Management interface. Specifically, it is the IP address assigned to the interface specified on the Security Management appliance in the Quarantine IP Interface setting on the Spam Quarantine Settings page under Management Appliance > Centralized Services > Spam Quarantine.

To view the IP address for the specified interface: On the Security Management appliance, go to Management Appliance > Network > IP Interfaces and click the interface name.

Step 6 Enter the port number to use for delivering spam and suspected spam. The default is 6025. The port number that you enter here must match the Quarantine Port number that you entered on the Spam Quarantine Settings page under Management Appliance > Centralized Services > Spam Quarantine.

Step 7 For simplicity, configure the safelist/blocklist feature later. For more information and complete instructions, see Configuring and Managing the End User Safelist/Blocklist Feature, page 7-11.

Step 8 Submit and commit your changes.

Step 9 Disable the local spam quarantine. See “Disabling the Local Spam Quarantine” in the *Cisco IronPort AsyncOS for Email Security Daily Management Guide.*

Ignore any warnings to adjust Mail Policies as a result of this change. Mail Policies will use the external spam quarantine.

Step 10 Repeat for each managed Email Security appliance.

Adding the Centralized Spam Quarantine 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.

Procedure

Step 1 On the Security Management appliance, 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 the Email Security appliance.

b. Select the **Spam Quarantine** 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 Cisco IronPort 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 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 Spam Quarantine.

Step 6 Commit your changes.

Configuring Administrative User Access to the Cisco IronPort Spam Quarantine

Use the procedure in this section to allow administrative users with the following roles to manage messages in the Cisco IronPort Spam Quarantine: Operator, Read-Only Operator, Help Desk, or Guest roles, and custom user roles that include access to the Spam Quarantine.

Administrator-level users, including the default admin user and Email Administrator users, can always access the Spam Quarantine and do not need to be associated with the Spam Quarantine feature using this procedure.

Note Non-Administrator-level users can access messages in the Spam Quarantine, but they cannot edit the quarantine settings. Administrator-level users can access messages and edit the settings.

To enable administrative users who do not have full Administrator privileges to manage messages in the Spam Quarantine:

Procedure

Step 1 Make sure you have created users and assigned them a user role with access to the Spam Quarantine. For information, see About Distributing Administrative Tasks, page 12-1.

Step 2 On the Security Management appliance, choose **Management Appliance > Centralized Services > Spam Quarantine**.

Step 3 Click **Enable** or **Edit Settings** in the Spam Quarantine Settings section.

Step 4 In the Administrative Users area of the Spam Quarantine Settings section, click the selection link for Local Users, Externally Authenticated Users, or Custom User Roles.

Step 5 Select the users to whom you want to grant access to view and manage messages in the Spam Quarantine.

Step 6 Click OK.

Step 7 Repeat if needed for each of the other types of Administrative Users listed in the section (Local Users, Externally Authenticated Users, or Custom User Roles).

Step 8 Submit and commit your changes.

Configuring Spam Management Features for End Users

- Configuring End User Quarantine Access, page 7-8.
- Configuring Spam Notifications for End Users, page 7-9.
- Configuring and Managing the End User Safelist/Blocklist Feature, page 7-11

Note You can configure one of the additional settings, but not the others. For example, to provide access only upon request or only to specified users, you might configure end user access, but not spam notifications.

Configuring End User Quarantine Access

You can allow email users to manage their own messages in the Cisco IronPort Spam Quarantine.

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > Centralized Services > Spam Quarantine**.

Step 2 Click the **Edit Settings** in the Cisco IronPort Spam Quarantine Settings section.

Step 3 Scroll down to the **End-User Quarantine Access** section.

Step 4 Check the **Enable End-User Quarantine Access** check box.

Step 5 Specify the method to authenticate end users when they attempt to view their quarantined messages. You can use mailbox authentication, LDAP authentication, or none.

- **Mailbox authentication:** For sites without an LDAP directory for authentication, the quarantine can validate users' email addresses and passwords against a standards-based IMAP or POP server that holds their mailboxes. When logging in to the web UI, users enter their full email address and mailbox password. The quarantine uses this information to log in to the mailbox server as the user. If the login is successful, the user is authenticated and the quarantine logs out of the mailbox server without making any changes to the user's Inbox. Mailbox authentication is recommended for sites that do not use an LDAP directory. However, mailbox authentication cannot provide a user with quarantined messages that were sent to multiple email aliases.

  Select the type of mailbox server (IMAP or POP). Specify a server name and whether or not to use SSL for a secure connection. Enter a port number for the server. Supply a domain (for example, company.com) to append to unqualified user names.

  If the POP server advertises APOP support in the banner, then for security reasons (namely, to avoid sending the password in the clear) the appliance uses APOP only. If APOP is not supported for some users, then the POP server should be reconfigured so that it does not advertise APOP.

  - **LDAP**: If you do not have an LDAP server or an active end user authentication query set up, choose **Management Appliance > System Administration > LDAP** to configure your LDAP server settings and end user authentication query string. For information about configuring LDAP authentication, see Creating the LDAP Server Profile, page 10-2.
  - **None:** You can allow end user access to the Cisco IronPort Spam Quarantine without enabling authentication. In this case, users can access the quarantine by clicking a link in the notification message, and the system does not use mailbox or LDAP authentication.

Step 6 Specify whether or not to display message bodies before messages are released from the quarantine.

If this check box is selected, users cannot view the message body through the Cisco IronPort Spam Quarantine page. Instead, to view a quarantined message, users must release the message and view it in their mail application (for example, Microsoft Outlook). You can use this feature for policy and regulation compliance — for example, if a regulation requires that all viewed email be archived.

Step 7 Submit and commit your changes.

Configuring Spam Notifications for End Users

Spam notifications are email messages sent to email users when they have messages in the Cisco IronPort Spam Quarantine. Notifications contain a list of quarantined spam or suspected spam for the user. Notifications also include a link for users to view their quarantined messages. Once enabled, notifications are sent according to the schedule that you specify.

Spam notifications can provide a way for end users to log in to the quarantine without using LDAP or mailbox authentication. Users access the quarantine through the email notifications that they receive (if notifications are enabled for the quarantine). Clicking a message subject logs the user in to the web UI for the quarantine.

Note This login method does not display quarantined messages for other aliases that the end user may have. Also, if the notification was sent to a distribution list that was expanded after the appliance processed it, then multiple recipients might have access to the same quarantine for the list.

Because of the way the appliance generates spam notifications, users may receive multiple spam notifications for their email aliases or if they use multiple email addresses. You can use the alias consolidation feature to prevent some occurrences of multiple notifications. **If you do not have an**

LDAP server or an active alias consolidation query set up, navigate to Management Appliance > **System Administration > LDAP to configure your LDAP server settings and alias consolidation query string.** For more information, see Configuring Spam Management Features for End Users, page 7-8.

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > Centralized Services > Spam Quarantine**.

Step 2 Click **Edit Settings** in the Cisco IronPort Spam Quarantine Settings section.

Step 3 Check the **Enable Spam Notification** check box to enable spam notification.

Step 4 Enter a **From Address** for the notifications. Users may want to add this address to a "whitelist" supported by their email client.

Step 5 Enter a **Subject** for the notification.

Step 6 Enter a customized **Title** for the notification.

Step 7 Select a **Default Language**.

Step 8 Customize the message body. AsyncOS supports several message variables that, when placed in the message body, are expanded to the actual value for the specific end user. For example, **%username%** is expanded to the actual user's name when the notification is generated for that user. The supported message variables include:

- **New Message Count (%new_message_count%):** number of new messages since the user's last login
- **Total Message Count (%total_message_count%):** number of messages for the user in the end user 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%):** list of new messages in the quarantine for the user

You can include these message variables in the message body by entering them directly in the text of the Message Body field, or you can place the cursor where you want the variable inserted and then click the name of the variable in the Message Variables list on the right.

Step 9 Select a message format (HTML, Text, or HTML/Text).

Step 10 Specify a bounce address. Bounced notifications are sent to this address.

Step 11 Optionally, you can consolidate messages sent to the same LDAP user at different addresses.

Step 12 Set the notification schedule. You can configure the notifications to be sent once a month, once a week, or at specified times during the day (weekdays only or including weekends).

Step 13 Submit commit your changes.

Configuring and Managing the End User Safelist/Blocklist Feature

You can allow end users to create safelists and blocklists to better control which email messages are treated as spam. Safelists allow a user to ensure that mail from specified users and domains is never treated as spam. Blocklists ensure that mail from other users and domains is always treated as spam. When you enable the safelist/blocklist feature, each end user can maintain a safelist and blocklist for his or her email account.

Note A safelist or blocklist setting does not prevent the Email Security appliance from scanning a message for viruses or determining if the message meets the criteria for a content-related mail policy. If a message is sent from a safelist member, it might not be delivered to the end user depending on other scanning settings.

When a user adds an entry to a safelist or blocklist, the entry is stored in a database on the Security Management appliance and periodically updated and synchronized on all related Email Security appliances. For information about synchronization, see Synchronizing Safelist and Blocklist Settings and Databases, page 7-12. For information on backing up the database, see Backing Up and Restoring the Safelist/Blocklist Database, page 7-13.

The safelists and blocklists are created and maintained by end users. However, an administrator enables the feature and configures delivery settings for email messages that match entries in the blocklist. Because the safelists and blocklists are related to the Cisco IronPort Spam Quarantine, delivery behavior is also contingent on other anti-spam settings. A message might skip anti-spam scanning based on the processing that occurs before the message reaches the Email Security Manager in the email pipeline. For more information about message processing, see "Understanding the Email Pipeline" in the *Cisco IronPort AsyncOS for Email Security Configuration Guide.*

For example, if you configure the "Accept" mail flow policy in the 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 mailflow policy that skips anti-spam scanning for certain message recipients, these recipients will not have their safelist and blocklist settings applied.

For more information about delivery of safelist/blocklist messages, see Message Delivery for Safelists and Blocklists, page 7-13.

Enabling and Configuring Safelist/Blocklists on the Security Management Appliance

Before you can enable the safelist/blocklist feature, you must enable the Cisco IronPort Spam Quarantine on the appliance. For more information about enabling the Cisco IronPort Spam Quarantine, see Setting Up the Centralized Spam Quarantine, page 7-2.

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > Centralized Services > Spam Quarantine**.

Step 2 Click **Enable** in the End-User Safelist/Blocklist section.

Step 3 Click **Edit Settings** in the End-User Safelist/Blocklist section.

Step 4 Verify that the **Enable End User Safelist/Blocklist Feature** check box is checked.

Step 5 Specify the maximum number of list items per user. This value is the maximum number of addresses and domains that a user can include in each safelist and blocklist. The default is 100.

Note If you allow a large number of list entries per user, system performance might be adversely affected.

Step 6 Select the update frequency. This value determines how often AsyncOS updates the safelist/blocklist databases on the Email Security appliances in the system. The default is every two hours for M10, M600, and M650 appliances. The default is every four hours for M1000 and M1050 appliances.

Step 7 Submit and commit your changes.

Step 8 Configure the Email Security appliance to support centralization of this feature. See Configuring Safelist/Blocklist Settings on the Email Security Appliance, page 7-12.

Configuring Safelist/Blocklist Settings on the Email Security Appliance

To configure the centralized safelist/blocklist settings on managed Email Security appliances, complete the following steps **on each** Email Security appliance:

Procedure

Step 1 On the Email Security appliance, navigate to **Security Services > External Spam Quarantine**.

Step 2 Click the **Edit Settings** button.

Step 3 Select the check box to enable the end user safelist/blocklist feature.

Step 4 Select whether to quarantine or delete messages from blocklisted senders.

Step 5 Submit and commit your changes.

Step 6 Repeat for each managed Email Security appliance.

Synchronizing Safelist and Blocklist Settings and Databases

With the Security Management appliance, you can easily synchronize the safelist/blocklist databases on all managed appliances.

Note Before you can synchronize safelist/blocklist databases, you need to enable the safelist/blocklist feature and add at least one managed appliance to the Security Management appliance. For more information about adding managed appliances, see About Adding Managed Appliances, page 2-12.

To synchronize safelist/blocklist databases, click the **Synchronize All Appliances** button on the Management Appliance > Centralized Services > Spam Quarantine page.

If you use the centralized management feature to configure multiple appliances, you can configure administrator settings using centralized management. If you do not use centralized management, you can manually verify that settings are consistent across machines.

For more information about using FTP to access appliances, see Appendix A, `IP Interfaces and Accessing the Appliance,' on page 1.

Message Delivery for Safelists and Blocklists

When you enable safelists and blocklists, the Email Security appliance scans the messages against the safelist/blocklist database immediately before anti-spam scanning. If the appliance detects a sender or domain that matches an end user's safelist/blocklist setting, the message is splintered if it has multiple recipients with different safelist/blocklist settings. For example, sender X sends a message to both recipient A and recipient B. Recipient A has safelisted sender X, but recipient B has no entry for the sender in either 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 it skips anti-spam scanning. The message bound for recipient B is scanned with the anti-spam scanning engine. Both messages then continue along the pipeline (through anti-virus scanning, content policies, and so forth), and they are subject to any configured settings.

If a message sender or domain is blocklisted, the delivery behavior depends on the blocklist action settings. 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.

Note You specify blocklist actions in the external spam quarantine settings on the Email Security appliance. For more information, see Configuring the Email Security Appliance for Centralized Spam Quarantine, page 7-5.

If you configure the blocklist action to quarantine messages, the message is scanned and eventually quarantined. If you configure the blocklist action to delete messages, the message is deleted immediately after safelist/blocklist scanning.

Backing Up and Restoring the Safelist/Blocklist Database

To maintain a backup of the safelist/blocklist database, the Security Management appliance enables you to save the database as a .csv file. The .csv file is maintained separately from the XML configuration file that contains the appliance configuration settings. If you upgrade your appliance or run the System Setup Wizard, first back up the safelist/blocklist database to the .csv file.

Note You can edit the .csv file and then upload it to modify individual end users' safelists and blocklists.

When you back up the database, the appliance saves the .csv file to the `/configuration` directory using the following naming convention:

slbl-<serial number>-<timestamp>.csv

When you back up your Security Management appliance, you can choose whether or not to include the Safelist/Blocklist database. See Backing Up Your Security Management Appliance, page 13-6.

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > System Administration > Configuration File**.

Step 2 Scroll down to the **End-User Safelist/Blocklist Database (Spam Quarantine)** section.

Step 3 Click **Backup Now** to back up the database to a .csv file.

Step 4 Click **Select File to Restore** to restore the database.

The appliance displays a list of backup files that are stored in the `/configuration` directory.

Step 5 Select the Safelist/Blocklist backup file that you want to restore, and click **Restore**.

Troubleshooting Safelists and Blocklists

End users maintain their own safelists and blocklists. Administrators can access an end user's safelist or blocklist by logging in to the end user account with the user's login and password. Alternatively, an administrator can download a backup version of the safelist/blocklist database to edit individual users' lists.

To troubleshoot issues with safelists and blocklists, you can view the log files or system alerts.

When an email message is blocked due to safelist/blocklist settings, the action is logged in the ISQ_logs or the anti-spam log files.

Alerts are sent out when the database is created and updated, or if there are errors in modifying the database or running the safelist/blocklist processes.

For more information about alerts, see Managing Alerts, page 13-28.

For more information about log files, see Chapter 14, "Logging.".

Using End User Safelists and Blocklists

End users can create safelists to ensure that messages from specified senders are never treated as spam, and they can use blocklists to ensure that messages from specified senders are always treated as spam. For example, an end user might receive unwanted email from a mailing list. The user can add this sender to the user's blocklist to prevent email messages from the sender from being delivered. On the other hand, an end user might find that email messages from a legitimate sender are sent to the Cisco IronPort Spam Quarantine, but those email messages should not be treated as spam. To prevent mail from that sender from being quarantined, the user can add the sender to the user's safelist.

Note Safelist/blocklist settings are contingent on other settings configured by the system administrator. For example, a safelisted message may not be delivered if it is determined to be virus positive, or if the administrator determines that the content does not conform to company email policies.

Accessing Safelists and Blocklists

To access safelists and blocklists, end users whose accounts are authenticated using LDAP or mailbox authentication (IMAP or POP) must log in to their accounts on the Cisco IronPort Spam Quarantine. The end users must log in to their accounts even if they are accustomed to accessing messages through a spam notification (which typically does not require LDAP or mailbox authentication). If the end user authentication is set to None, end users do not need to log in to their accounts to access safelist/blocklist settings.

Adding Entries to Safelists and Blocklists

Entries (including those using IPv6 addresses) can be added to safelists and blocklists using the following formats:

- user@domain.com

  user@[203.0.113.15]

  user@[ipv6:2001:db8:80:1::5]
- server.domain.com
- domain.com

  [203.0.113.15]

  [ipv6:2001:db8:80:1::5]

End users cannot add a sender or domain to both their safelist and their blocklist at the same time. However, if they add a domain to a safelist and a user in that domain to the blocklist (or vice versa), the appliance applies both rules. For example, if the end user adds *example.com* to the safelist, and adds *george@example.com* to the blocklist, the appliance delivers all mail from example.com without scanning for spam, but it treats mail from *george@example.com* as spam.

End users cannot allow or block a range of subdomains using the following syntax: .domain.com. However, an end user can explicitly block a specific domain using the following syntax: server.domain.com.

Working with Safelists

End users can add senders to safelists in two ways. From the Cisco IronPort Spam Quarantine, they can manually add a sender to the safelist by clicking the Options menu in the upper-right corner of the GUI and then selecting Safelist.

Figure 7-1 Safelist in End User Quarantine

Add an email address or domain to the list, and click **Add to List.**

End users can also add senders to the safelist if the message has been sent to the Cisco IronPort Spam Quarantine. If the message from a particular sender is held in the Cisco IronPort Spam Quarantine, the end user can select the check box next to the message, and choose “Release and Add to Safelist” from the drop-down menu.

Figure 7-2 Safelist in End User Quarantine

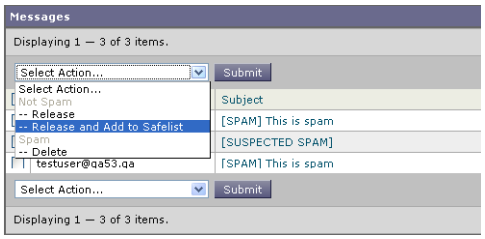

The envelope sender and the From header for the specified mail are both added to the safelist, and the released messages proceed directly to the destination queue, skipping any further work queue processing in the email pipeline.

Note End users can also use the spam notification message to release messages. Click the Not Spam link to release a particular message. End users also have the option to add senders to their safelists.

Working with Blocklists

End users can use blocklists to prevent the delivery of mail from specified senders. To add senders to a blocklist, the end user selects Options > Blocklist from the end user quarantine.

Figure 7-3 Adding Senders to a Blocklist

From the end user quarantine, the end user enters an email address or domain in the field, and clicks **Add to List.**

When the Email Security appliance receives mail from the specified email address or domain that matches an entry in the blocklist, it treats the mail as spam. The mail might be deleted or quarantined, depending on the blocklist action setting.

Managing Messages in the Cisco IronPort Spam Quarantine

This section explains how administrators can manage messages in the Cisco IronPort Spam Quarantine. When an administrator views the quarantine, all of the messages contained in the quarantine are available.

Note The GUI for viewing and managing messages is slightly different for end users who access the Cisco IronPort Spam Quarantine. For information about the end user GUI, access the Cisco IronPort Spam Quarantine as an end user, and view the online help.

As an administrator, you can perform the following actions on messages in the Cisco IronPort Spam Quarantine:

- View messages
- Deliver messages
- Delete messages
- Search messages

Searching for Messages in the Cisco IronPort Spam Quarantine

Procedure

Step 1 On the Security Management appliance, choose **Email > Message Quarantine > Spam Quarantine**.

Step 2 Click the **Spam Quarantine** link.

Step 3 On the search form, enter the dates to search. You can search messages from the current day or the past week, or you can click the calendar icons to select a date range.

Step 4 Optionally, specify a text string for the From address, To address, or message subject. Select whether the search results contain, do not contain, match exactly, start with, or end with the value that you enter.

Step 5 Optionally, specify an envelope recipient. Select whether the search results contain, do not contain, match exactly, start with, or end with the value that you enter.

The envelope recipient is the address of the email message recipient as defined in the "RCPT TO" SMTP command. The envelope recipient is also sometimes called the "Recipient To" address or the "Envelope To" address.

Step 6 Click **Search.**

Messages that match the search criteria appear below the Search section of the page.

Searching Large Message Collections

If a large number of messages are stored in the Cisco IronPort Spam Quarantine and the search terms are not narrowly defined, a query may take a long time to display the search results, or the query may time out.

You are prompted to confirm whether you want to resubmit the search.

Note Running multiple large searches simultaneously can adversely affect the performance of the appliance.

Viewing Messages in the Cisco IronPort Spam Quarantine

The message list shows messages in the Cisco IronPort Spam Quarantine. You can select how many messages appear on a page. You can sort the display by clicking the column headings. Click a column heading a second time to reverse the sorting.

Click the subject of a message to view the message, including the body and headers. The first 20K of the message appears on the Message Details page. If the message is longer, it is truncated at 20K. Click the link at the bottom of the page to view the rest of the message.

From the Message Details page, you can select Delete to delete a message or select Release to release the message from the quarantine. Releasing a message causes it to be delivered.

Viewing HTML Messages

The Cisco IronPort Spam Quarantine attempts to render an approximation of HTML-based messages. Images do not appear.

Viewing Encoded Messages

Base64 encoded messages are decoded and then appear.

Delivering Messages in the Cisco IronPort Spam Quarantine

To release messages for delivery, select the check box next to the messages and click **Release.**

Select the check box in the heading row to select all of the messages 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 Cisco IronPort Spam Quarantine

The Cisco IronPort Spam Quarantine can be configured to automatically delete messages after a specified period of time. You can also delete messages from the Cisco IronPort Spam Quarantine manually.

To delete specific messages, select the check box next to the messages you want to delete and then click **Delete.** Select the check box in the heading row to select all of the messages displayed on the page.

To delete all of the messages in the Cisco IronPort Spam Quarantine, disable the quarantine from the **Management Appliance > Centralized Services > Spam Quarantine** page, and then click the **Delete All** link that appears.

CHAPTER 8

Managing Web Security Appliances

- About Centralized Configuration Management, page 8-1
- Determining the Correct Configuration Publishing Method, page 8-1
- Setting Up Configuration Masters, page 8-2
- Setting Up to Use Advanced File Publishing, page 8-12
- Publishing Configurations to Web Security Appliances, page 8-13
- Viewing Status and History of Publishing Jobs, page 8-18
- Viewing Web Security Appliance Status, page 8-19
- URL Category Set Updates and Centralized Configuration Management, page 8-22

About Centralized Configuration Management

Centralized configuration management allows you to publish configurations from a Security Management appliance to associated Web Security appliances, in order to:

- Simplify and speed management of web security policies by configuring or updating settings once on the Security Management appliance, instead of on each Web Security appliance.
- Ensure uniform policy enforcement across distributed networks.

There are two ways to publish settings to Web Security appliances:

- Using Configuration Masters
- Using configuration files from Web Security appliances (using Advanced File Publishing)

Determining the Correct Configuration Publishing Method

There are two different processes for publishing configurations from the Security Management appliance, and each publishes different settings. Some settings cannot be centrally managed.

Generally:

- Use a Configuration Master for:

  Features that appear under the Web Security Manager menu on the Web Security appliance, such as policies and custom URL categories.

Exception: L4 Traffic Monitor (L4TM) settings are not included in Configuration Masters.

The exact features supported depend on the configuration master version, which corresponds to an AsyncOS for Web Security version.

Some features that are configurable in a Configuration Master also require configurations directly on the Web Security appliance in order to work. For example, SOCKS Policies are configurable via Configuration Master, but a SOCKS Proxy must first be configured directly on the Web Security appliance.

- Use a Configuration File (Advanced File Publishing) for:

  Features related to managing the appliance, for example configuring log subscriptions or alerts, or distributing administrative responsibilities.

  Exceptions:

  You cannot use a Security Management appliance to enable or configure the following features and functionality on Web Security appliances: FIPS mode for Federal Information Processing Standard, Network/interface settings, DNS, Web Cache Communication Protocol (WCCP), upstream proxy groups, certificates, the proxy mode, time settings such as SNTP, Level 4 Traffic Monitor (L4TM) settings, and authentication redirect hostname.

  You must configure these settings directly on your managed Web Security appliances. See the *Cisco IronPort AsyncOS for Web Security User Guide.*

Setting Up Configuration Masters

To set up for centralized configuration management using Configuration Masters, follow the procedures in Overview of Setting Up Configuration Masters, page 8-2, in order.

To prepare to use Advanced File Publishing only, see Setting Up to Use Advanced File Publishing, page 8-12 instead.

Overview of Setting Up Configuration Masters

To set up your system to centrally manage your Web Security appliances, follow these steps in order:

Step 1 **(Optional) Web Security appliance.** If you have a working Web Security appliance that can serve as a configuration model for all of your Web Security appliances, download a configuration file from that Web Security appliance. You can use this file to speed configuration of a Configuration Master in the Security Management appliance. For instructions, see "Saving and Loading the Appliance Configuration" in the *Cisco IronPort AsyncOS for Web Security User Guide.*

For compatibility of configuration files and Configuration Master versions, see the Release Notes for this release at http://www.cisco.com/en/US/products/ps10155/prod_release_notes_list.html.

Step 2 Check for general configuration requirements and caveats. See Important Notes About Using Configuration Masters, page 8-3.

Step 3 Determine the Configuration Master version to use for each Web Security appliance. See Determine the Configuration Master Version(s) to Use, page 8-3.

Step 4 **Security Management appliance.** Enable and configure the Cisco IronPort Centralized Configuration Manager. See Enabling Centralized Configuration Management on the Security Management Appliance, page 8-4.

Step 5 **Security Management appliance.** Initialize the Configuration Masters. See Initializing Configuration Masters, page 8-4.

Step 6 **Security Management appliance.** Associate Web Security appliances to the Configuration Masters. See About Associating Web Security Appliances to Configuration Masters, page 8-5.

Step 7 **Security Management appliance.** Import and/or manually configure policies, custom URL categories, and/or a web proxy bypass list in the Configuration Masters. See Configuring Settings to Publish, page 8-6.

Step 8 **Security Management appliance.** Ensure that the features enabled on each Web Security appliance match the features enabled for the Configuration Master assigned to that appliance. See Ensuring that Features are Enabled Consistently, page 8-10.

Step 9 **Security Management appliance.** After you have set up required Configuration Masters and enabled appropriate features, publish configurations to your Web Security appliances. See Publishing a Configuration Master, page 8-13.

Important Notes About Using Configuration Masters

Note On each Web Security appliance that you will manage centrally, check to be sure that all Realm Names in Network > Authentication are unique across appliances, unless the settings for same-name realms are identical.

Determine the Configuration Master Version(s) to Use

AsyncOS for Security Management provides multiple Configuration Masters, making it possible for you to centrally manage a heterogeneous deployment in which Web Security appliances run different versions of AsyncOS for Web Security that support different features.

Each Configuration Master contains configurations to be used for a particular version of AsyncOS for Web Security.

To determine which configuration master version(s) to use for your version(s) of AsyncOS for Web Security, see the SMA Compatibility Matrix, page 2-2.

Note For best results, the Configuration Master version should match the AsyncOS version on the Web Security appliance. Publishing an older Configuration Master version to a newer Web Security appliance may fail if settings on the Web Security appliance do not match the settings in the Configuration Master. This can occur even if the Web Appliance Status detail page does not indicate any discrepancies. In this case, you would need to manually compare the configurations on each appliance.

Enabling Centralized Configuration Management on the Security Management Appliance

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > Centralized Services > Web > Centralized Configuration Manager**.

Step 2 On the Centralized Configuration Manager page, click **Enable.**

Step 3 If you are enabling the Cisco IronPort Centralized Configuration Manager 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.

Initializing Configuration Masters

Note Before you copy or import Configuration Master 7.1 into Configuration Master 7.5 or 7.7, see Be Aware: Before You Set Up Configuration Master 7.5 and 7.7, page 8-23.

Procedure

Step 1 On the main Security Management appliance, choose **Web > Utilities > Configuration Masters**

Step 2 Click **Initialize** in the Options column.

Step 3 On the Initialize Configuration Master page:

- If you have an existing Configuration Master for a previous release and you want to use or start with the same settings for the new Configuration Master, choose **Copy Configuration Master**.

  You can also import settings from an existing Configuration Master later.
- Otherwise, choose **Use default settings**.

Step 4 Click **Initialize**.

The Configuration Master is now available.

Step 5 Repeat for each Configuration Master version to initialize.

Note After you initialize a configuration master, the Initialize option is not available. Instead, populate the Configuration Master using one of the methods described in Configuring Settings to Publish, page 8-6.

About Associating Web Security Appliances to Configuration Masters

For each Web Security appliance that you want to centrally manage, the policy configuration must be associated to the Configuration Master that matches the appliance's AsyncOS version. For example, if the Web Security appliance is running AsyncOS 7.7 for Web, then it should be associated to Configuration Master 7.7.

The simplest process for doing this depends on the situation:

If	Use This Procedure
You have not yet added Web Security appliances to the Security Management appliance	Adding Web Security Appliances and Associating Them with Configuration Master Versions, page 8-5
You have already added Web Security appliances	Associating Configuration Master Versions to Web Security Appliances, page 8-6

Adding Web Security Appliances and Associating Them with Configuration Master Versions

Use this procedure if you have not yet added your Web Security appliances to be centrally managed.

Procedure

Step 1 Verify compatibility of your Web Security appliances with available Configuration Master versions in this release of AsyncOS for Security Management. See the SMA Compatibility Matrix, page 2-2.

Step 2 On the Security Management appliance, choose **Management Appliance > Centralized Services > Security Appliances**.

Step 3 Click **Add Web Appliance**.

Step 4 In the Appliance Name and IP Address text fields, type the appliance name and the IP address for the Management interface of the Cisco IronPort 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**.

Step 5 The Centralized Configuration Manager service is pre-selected.

Step 6 Click **Establish Connection.**

Step 7 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.

Step 8 Wait for the Success message to appear above the table on the page.

Step 9 Choose the Configuration Master version to which you want to assign the appliance.

Step 10 Submit and commit your changes.

Step 11 Repeat this procedure for each Web Security Appliance for which you want to enable Centralized Configuration Management.

Associating Configuration Master Versions to Web Security Appliances

If you have already added Web Security appliances to the Security Management appliance, you can use the following procedure to quickly associate Web Security appliances with Configuration Master versions:

Procedure

Step 1 Verify compatibility of your Web Security appliances with available Configuration Master versions in this release of AsyncOS for Security Management. See the SMA Compatibility Matrix, page 2-2.

Step 2 On the Security Management appliance, choose **Web > Utilities > Configuration Masters**.

Note If a Configuration Master shows as Disabled, you can enable it by clicking Web > Utilities > Security Services Display, then **Edit Display Settings**. Select the check box for that Configuration Master to enable it. For more information, see Enabling Features to Publish, page 8-11.

Step 3 Click **Edit Appliance Assignment List**.

Step 4 In the rows of the appliances you want to associate, click to enter check marks in the **Masters** columns.

Step 5 Submit and commit your changes.

Configuring Settings to Publish

Set up your configuration masters with the settings you want to publish.

There are several ways to set up Configuration Masters:

- If you are upgrading from a previous release of AsyncOS for Security Management: If you did not initialize a new Configuration Master version by copying an earlier, existing Configuration Master into the new version, you can import the old version. See Importing from an Existing Configuration Master, page 8-7.
- If you have already configured a Web Security appliance and want to use those same configurations for multiple Web Security appliances: Import a saved configuration file from that appliance into the Configuration Master. (You may have saved this configuration file when you reviewed Setting Up Configuration Masters, page 8-2.)

  To import, see Importing Settings from a Web Security Appliance, page 8-7.
- If you want to modify imported settings, see Configuring Web Security Features Directly in Configuration Masters, page 8-8.
- If you have not yet configured policy settings, URL categories, or bypass settings on a Web Security appliance, you can configure these settings in an appropriate Configuration Master on the Security Management appliance.

For more information, see Configuring Web Security Features Directly in Configuration Masters, page 8-8.

Importing from an Existing Configuration Master

You can upgrade an existing configuration master to a new configuration master version. For example, you can import your Configuration Master 7.1 settings into Configuration Master 7.5 or 7.7.

Note Before you copy or import to Configuration Master 7.5 or 7.7, see Be Aware: Before You Set Up Configuration Master 7.5 or 7.7, page 8-23.

Procedure

Step 1 On the Security Management appliance, choose **Web > Utilities > Configuration Masters**.

Step 2 In the Options column, click **Import Configuration**.

Step 3 For **Select Configuration Source**, select a Configuration Master from the list.

Step 4 Choose whether or not to include existing custom user roles in this configuration.

For more information about custom user roles, see About Custom Web User Roles, page 12-7.

Step 5 Click **Import**.

Importing Settings from a Web Security Appliance

If you want to use an existing, working configuration from one of your Web Security appliances, you can import the configuration file to the Security Management appliance to create default policy settings for a Configuration Master.

For compatibility of configuration files and Configuration Master versions, see the Release Notes for this release at http://www.cisco.com/en/US/products/ps10155/prod_release_notes_list.html.

Caution You can import compatible web configuration files as often as you want, even if you have already published configurations to your managed Web Security appliances. However, be aware that importing a configuration file to a Configuration Master completely overwrites the settings associated with the selected Configuration Master. In addition, the security services settings on the Security Services Display page are set to match the imported configuration.

To populate the Configuration Master with a web configuration file:

Procedure

Step 1 Save a configuration file from the Web Security appliance.

Step 2 On the Security Management appliance, choose **Web > Utilities > Configuration Masters**.

Step 3 In the Options column, click **Import Configuration**.

Step 4 From the Select Configuration drop-down list, select **Web Configuration File**.

Step 5 In the New Master Defaults section, click **Browse** and select a valid configuration file from a Web Security appliance.

Step 6 Click **Import File.**

Step 7 Click **Import.**

Configuring Web Security Features Directly in Configuration Masters

You can configure the following features in a Configuration Master, depending on the version:

Configuration Master 7.1	Configuration Master 7.5	Configuration Master 7.7
• Identities • SaaS Policies • Decryption Policies • Routing Policies • Access Policies • Overall Bandwidth Limits • Cisco IronPort Data Security • External Data Loss Prevention • Outbound Malware Scanning • Custom URL Categories • Defined Time Ranges • Bypass Settings	• Identities • SaaS Policies • Decryption Policies • Routing Policies • Access Policies • Overall Bandwidth Limits • Cisco IronPort Data Security • External Data Loss Prevention • Outbound Malware Scanning • Custom URL Categories • Defined Time Ranges • Bypass Settings	• Identities • SaaS Policies • Decryption Policies • Routing Policies • Access Policies • Overall Bandwidth Limits • Cisco IronPort Data Security • External Data Loss Prevention • Outbound Malware Scanning • SOCKS Policies • Custom URL Categories • Defined Time Ranges • Bypass Settings

To configure settings for each feature directly in a configuration master, choose **Web > Configuration Master** <*version*> **>** <*feature*>.

Except for the few items described in SMA-Specific Differences when Configuring Features in Configuration Masters, page 8-9, instructions for configuring features in a Configuration Master are the same as instructions for configuring the same features on the Web Security appliance. For instructions, see the online help in your Web Security appliance or the *Cisco IronPort AsyncOS for Web Security User Guide* for the AsyncOS version corresponding to the Configuration Master version. If necessary, consult the following topic to determine the correct Configuration Master for your Web Security appliance: Determine the Configuration Master Version(s) to Use, page 8-3.

All versions of Web Security user guides are available from http://www.cisco.com/en/US/products/ps10164/products_user_guide_list.html.

SMA-Specific Differences when Configuring Features in Configuration Masters

When you configure a feature in a Configuration Master, note the following differences from configuring the same feature directly on the Web Security appliance.

Table 8-1 Feature Configuration: Differences between Configuration Master and Web Security Appliance

Feature or Page	Details
All features, especially new features in each release	For each feature that you configure in a Configuration Master, you must enable the feature in the Security Management appliance under Web > Utilities > Security Services Display. For more information, see Ensuring that Features are Enabled Consistently, page 8-10.
Identities	• See Tip for Working with Identities in Configuration Masters, page 8-9. • If you have realms on different Web Security appliances that have the same name but different protocols, choose the appropriate scheme for each desired realm in the Configuration Master. • The **Identify Users Transparently** option when adding or editing an Identity is available when a Web Security appliance with an authentication realm that supports transparent user identification has been added as a managed appliance. This feature was introduced in Configuration Master 7.5.
SaaS Policies	The authentication option “Prompt SaaS users who have been discovered by transparent user identification” is available only when a Web Security appliance with an authentication realm that supports transparent user identification has been added as a managed appliance.
Access Policies > Edit Group	When you configure the Identities and Users option in the Policy Member Definition section, the following applies if you use external directory servers: When you search for groups on the Edit Group page, only the first 500 matching results are displayed. If you do not see the desired group, you can add it to the “Authorized Groups” list by entering it in the Directory search field and clicking the "Add" button.
Access Policies > Web Reputation and Anti-Malware Settings	Options available on this page depend on whether Adaptive Scanning is enabled for the relevant configuration master. Check this setting in Web > Utilities > Security Services Display. This feature was introduced in Configuration Master 7.5.

Tip for Working with Identities in Configuration Masters

When creating an Identity on the Security Management appliance, you have the option of making it apply only to specific appliances. So for example, if you purchase a Security Management appliance and want to preserve the existing Web Security appliance configurations and the policies that were created for each Web Security appliance, you must load one file into the machine, and then add policies from other machines by hand.

One way to accomplish this is to make a set of Identities for each appliance, then have policies which refer to those Identities. When the Security Management appliance publishes the configuration, those Identities and the policies which refer to them will automatically be removed and disabled. Using this method, you do not have to configure anything manually. This is essentially a ‘per-appliance’ identity.

The only challenge with this method is if you have a default policy or Identity that differs between sites. For example, if you have a policy set for “default allow with auth” at one site and a “default deny” at another. At this point you will need to create per-appliance Identities and policies just above the default; essentially creating your own “default” policy.

Ensuring that Features are Enabled Consistently

Before you publish a Configuration Master, you should ensure that it will publish and that the intended features will be enabled and configured as you expect them to be after publishing.

To do this, do both of the following:

- Comparing Enabled Features, page 8-10
- Enabling Features to Publish, page 8-11

Note If multiple Web Security appliances with different features enabled are assigned to the same Configuration Master, you should publish to each appliance separately, and perform these procedures before each publish.

Comparing Enabled Features

Verify that the features enabled on each Web Security appliance match the features enabled for the Configuration Master associated with that appliance.

Note If multiple Web Security appliances with different features enabled are assigned to the same Configuration Master, you should publish to each appliance separately, and perform this check before each publish.

To verify enabled features for a Web Security appliance:

Procedure

Step 1 On the Security Management appliance, choose **Web > Utilities > Web Appliance Status**.

Step 2 Click the name of a Web Security appliance to which you will publish a Configuration Master.

Step 3 Scroll to the **Security Services** table.

Step 4 Verify that the Feature Keys for all enabled features are active and not expired.

Step 5 Compare the settings in the **Services** columns:

The **Web Appliance Service** column and the **Is Service Displayed on Management Appliance?** column should be consistent.

- Enabled = Yes
- Disabled and Not Configured = No or Disabled.
- N/A = Not Applicable. For example, the option may not be configurable using a Configuration Master, but is listed so that you can see the Feature Key status.

Configuration mismatches will appear in red text.

Step 6 If the enabled/disabled settings for a feature do not match, do one of the following:

- Change the relevant setting for the Configuration Master. See Enabling Features to Publish, page 8-11.
- Enable or disable the feature on the Web Security Appliance. Some changes may impact multiple features. See the information about the relevant feature in the *Cisco IronPort AsyncOS for Web Security User Guide*.

Enabling Features to Publish

Enable the features for which you want to publish settings using a Configuration Master.

Note Enabling a feature for a Configuration Master does not enable the feature on the Web Security appliances.

Features enabled for each Configuration Master are summarized on the Security Services Display page. "N/A" alongside a feature indicates that the feature is not available in that Configuration Master version.

To enable features to publish:

Procedure

Step 1 Determine which features must be enabled and disabled. See Comparing Enabled Features, page 8-10.

Step 2 On the Security Management appliance, choose **Web > Utilities > Security Services Display**.

Step 3 Click **Edit Settings**.

The Edit Security Services Display page lists the features that appear in each Configuration Master.

Note Web Proxy is not listed as a feature, because it is assumed that the Web Proxy is enabled in order to execute any of the managed policy types on the Web Security appliances. If the Web Proxy is disabled, any policies published to the Web Security appliances will be ignored.

Step 4 (Optional) Hide Configuration Masters that you will not use. To avoid unintended effects, see the Note in Disabling Unused Configuration Masters, page 8-12.

Step 5 For each Configuration Master that you will use, select or uncheck the **Yes** check box for each feature to enable.

Special notes for certain features (available options vary by Configuration Master version):

- Transparent Mode. If you use Forward mode, the proxy bypass feature will not be available.
- HTTPS Proxy. HTTPS proxy must be enabled in order to configure decryption policies.
- Upstream Proxy Groups. Upstream proxy groups must be available on your Web Security appliances if you want to use routing policies.

Step 6 Make changes for each Configuration Master that you will use.

Step 7 Click **Submit.** The GUI displays specific warning messages if the changes you made to the security services settings will affect policies configured on your Web Security appliances. If you are sure that you want to submit your changes, click **Continue.**

Step 8 On the **Security Services Display** page, confirm that **Yes** appears alongside each option that you selected.

Step 9 Commit your changes.

Step 10 Verify that all features are now correctly enabled or disabled for the appliance(s) that you will publish to. See Comparing Enabled Features, page 8-10.

Disabling Unused Configuration Masters

You can choose not to display unused Configuration Masters.

As an example, the Configuration Master tabs and **Security Services Display** page look like the following when some Configuration Masters are disabled:

Note When a Configuration Master is disabled, all references to it are removed from the GUI including the corresponding Configuration Master tab. Pending publish jobs that use the Configuration Master are deleted, and all Web Security appliances assigned to the hidden Configuration Master are re-categorized as not assigned. At least one Configuration Master must be enabled.

Procedure

Step 1 On the Security Management appliance, choose **Web > Utilities > Security Services Display**.

Step 2 Click **Edit Settings**.

Step 3 Uncheck the checkbox(es) for unused Configuration Masters.

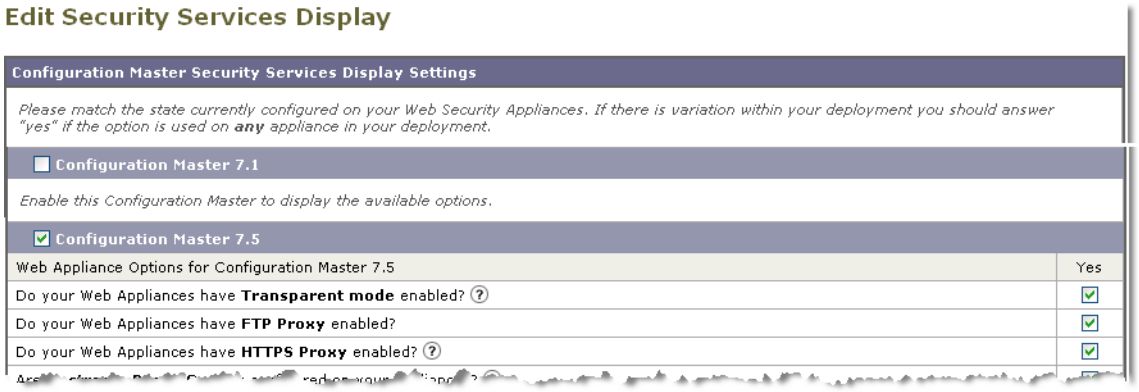

Step 4 Submit and commit your changes.

Setting Up to Use Advanced File Publishing

If your system is set up to use Configuration Masters, it is already set up for Advanced File Publishing.

Otherwise, complete procedures in the following topics, which apply to Advanced File Publishing as well as to publishing Configuration Masters.

- Enabling Centralized Configuration Management on the Security Management Appliance, page 8-4

- Initializing Configuration Masters, page 8-4
- About Associating Web Security Appliances to Configuration Masters, page 8-5

Publishing Configurations to Web Security Appliances

- Publishing a Configuration Master, page 8-13
- Publishing Configurations Using Advanced File Publishing, page 8-16

Publishing a Configuration Master

After editing or importing settings in a Configuration Master, you can publish them to the Web Security appliances associated with the Configuration Master.

- Before You Publish a Configuration Master, page 8-13
- Publishing a Configuration Master Now, page 8-14
- Publishing a Configuration Master Later, page 8-15
- Publishing a Configuration Master Using the Command Line Interface, page 8-16

Before You Publish a Configuration Master

Publishing a Configuration Master overwrites existing policy information on the Web Security appliances associated to that Configuration Master.

For information about which settings you can configure using a Configuration Master, see Determining the Correct Configuration Publishing Method, page 8-1.

Before you can publish a Configuration Master:

- The AsyncOS version on the target Web Security appliance must be the same as or newer than the Configuration Master version. For specific requirements, see the SMA Compatibility Matrix, page 2-2. We strongly recommend using Configuration Master 7.5 to publish to Web Security appliances running AsyncOS 7.5.
- (First time only) You must follow the procedures in Setting Up Configuration Masters, page 8-2.
- Save a configuration file from each target Web Security appliance so that you can restore the existing configuration in case of problems with the published configuration. See the *Cisco IronPort AsyncOS for Web Security User Guide* for details.
- To ensure that the Configuration Master will publish and that the intended set of features will be enabled after publishing, verify the feature sets of each Web Security appliance and the associated Configuration Master and make any needed changes. See Comparing Enabled Features, page 8-10 and if necessary, Enabling Features to Publish, page 8-11.

  If different features are enabled on Web Security appliances assigned to the same Configuration Master, you must publish to each appliance separately, and verify and enable features before each publish.
- If you have reverted AsyncOS on the target Web Security appliance, you may need to associate a different Configuration Master with that appliance.

- If you publish a Configuration Master to a Web Security appliance that does not have a realm configured with Transparent User Identification enabled, but you have selected Transparent User Identification in an Identity or SaaS Policy:
  - For Identities, Transparent User Identification is disabled and the Require Authentication option is selected instead.
  - For Saas Policies, the Transparent User Identification option is disabled and the default option (Always prompt SaaS users for proxy authentication) is selected instead.
- Any change that would cause a Web proxy restart when committed on the Web Security appliance will also cause a proxy restart when you publish it from the Security Management appliance. You will receive a warning in these situations.

  Proxy restarts may also occur on publish if a change requiring proxy restart has been made on the Web Security appliance. For example, if new groups are added on the Web Security appliance to a group authentication configuration for an access policy, the web proxy will restart the next time the configuration master is published. You will not receive warnings about proxy restarts in these cases.

  Web Proxy restarts temporarily interrupt web security services. For information about the effects of restarting the web proxy, see the "Checking for Web Proxy Restart on Commit" section in the *Cisco IronPort AsyncOS for Web Security User Guide*.
- When you publish any change to an Identity, all end-users must re-authenticate.

Note Publishing External DLP policies from a Security Management appliance to multiple Web Security appliances that are not configured for RSA servers is not an issue. When you try to publish, the Security Management appliance, will send the following publish status warning, "**The Security Services display settings configured for Configuration Master 7.1 do not currently reflect the state of one or more Security Services on Web Appliances associated with this publish request. The affected appliances are: "**[*WSA Appliance Name*]**". This may indicate a misconfiguration of the Security Services display settings for this particular Configuration Master. Go to the Web Appliance Status page for each appliance provides a detailed view to troubleshooting this issue. Do you want to continue publishing the configuration now?"**

If you decide to continue to publish, the Web Security appliance that is not configured for the RSA servers will receive the External DLP policies, but these policies will be disabled.The Web Security appliance External DLP page will not show the published policies if External DLP Server is not configured.

Publishing a Configuration Master Now

Procedure

Step 1 See important requirements and information in Before You Publish a Configuration Master, page 8-13.

Step 2 On the Security Management appliance, choose **Web > Utilities > Publish to Web Appliances.**

Step 3 Click **Publish Configuration Now.**

Step 4 "System-generated job name" is selected by default, or enter a user-defined job name (80 characters or fewer).

Step 5 Select the Configuration Master to publish.

Step 6 Select the Web Security appliances to which you want to publish the Configuration Master. Choose "All assigned appliances" to publish the configuration to all appliances assigned to the Configuration Master.

or

Choose "Select appliances in list" to display the list of appliances assigned to the Configuration Master. Select the appliances to which you want to publish the configuration.

Step 7 Click **Publish.**

Red progress bars and text on the Publish in Progress page indicate that an error occurred during publishing. If another job is currently publishing, then your request will be executed when the previous job is complete.

Note Details of the job in progress also appear on the **Web > Utilities > Publish to Web Appliances** page. Click **Check Progress** to access the Publish in Progress page.

Step 8 Check to be sure your publish was completely successful. See Viewing Publish History, page 8-18. Items that were not published completely will be noted.

Publishing a Configuration Master Later

Procedure

Step 1 See important requirements and information in Before You Publish a Configuration Master, page 8-13.

Step 2 On the Security Management appliance, choose **Web > Utilities > Publish to Web Appliances.**

Step 3 Click **Schedule a Job.**

Step 4 "System-generated job name" is selected by default, or enter a user-defined job name (80 characters or fewer).

Step 5 Enter the date and time when you want to publish the Configuration Master.

Step 6 Select the Configuration Master to publish.

Step 7 Select the Web Security appliances to which you want to publish the Configuration Master. Choose "All assigned appliances" to publish the configuration to all appliances assigned to the Configuration Master.

or

Choose "Select appliances in list" to display the list of appliances assigned to the Configuration Master. Select the appliances to which you want to publish the configuration.

Step 8 Click **Submit.**

Step 9 View a list of scheduled jobs on the **Web > Utilities > Publish to Web Appliances** page. To edit a scheduled job, click the name of the job. To cancel a pending job, click the corresponding trash can icon and confirm that you want to delete the job.

Step 10 You may want to create a reminder for yourself (for example, in your calendar) to check to be sure your publish was completely successful. See Viewing Publish History, page 8-18. Items that were not published completely will be noted.

Note If you reboot or upgrade the appliance before the scheduled publishing job occurs, you must reschedule the job.

Publishing a Configuration Master Using the Command Line Interface

Note See important requirements and information in Before You Publish a Configuration Master, page 8-13.

The Security Management appliance provides you with the ability to publish changes through a Configuration Master using the following CLI command:

publishconfig config_master [*--job_name*] [*--host_list* | *host_ip*]

where **config_master** is either 7.1, 7.5, or 7.7. This keyword is required. The option *job_name* is optional and will be generated if it is not specified.

The option *host_list* is a list of host names or IP addresses for Web Security appliances to be published, and will be published to all hosts assigned to the Configuration Master if not specified. The option *host_ip* can be multiple host IP addresses, each separated by a comma.

To verify that the **publishconfig** command was successful, check the **smad_logs** file. You can also verify that the publish history was successful from the Security Management appliance GUI by choosing **Web > Utilities > Web Appliance Status**. From this page choose the web appliance that you want the publish history details. Additionally, you can go the Publish History page: **Web > Utilities > Publish > Publish History**.

Publishing Configurations Using Advanced File Publishing

Use advanced file publish to push a compatible XML configuration file from your local file system to managed Web Security appliances.

For information about which settings you can configure using Advanced File Publishing, see Determining the Correct Configuration Publishing Method, page 8-1.

Note Any change that would cause a Web proxy restart when committed on the Web Security appliance will also cause a proxy restart when you publish it from the Security Management appliance. You will receive a warning about proxy restarts when you use Advanced File Publishing.

Web Proxy restarts temporarily interrupt web security services. For more information about the effects of restarting the web proxy, see the "Checking for Web Proxy Restart on Commit" section in the *Cisco IronPort AsyncOS for Web Security User Guide*.

To perform an advanced file publish, you can choose one of the following:

- Advanced File Publish: Publish Configuration Now, page 8-16
- Advanced File Publish: Publish Later, page 8-17

Advanced File Publish: Publish Configuration Now

Procedure

Step 1 From the source Web Security appliance, save a Configuration File.

For configuration file compatibility information, see the SMA Compatibility Matrix, page 2-2.

For instructions on saving a configuration file from a Web Security appliance, see the *Cisco IronPort AsyncOS for Web Security User Guide*.

Step 2 On each destination Web Security appliance, save the existing configuration on your Web Security appliance to a configuration file. See the *Cisco IronPort AsyncOS for Web Security User Guide* for details.

Step 3 On the main Security Management appliance window, choose **Web > Utilities > Publish to Web Appliances**.

Step 4 Click **Publish Configuration Now.**

Step 5 "System-generated job name" is selected by default, or enter a job name (up to 80 characters).

Step 6 For **Configuration Master to Publish**, select **Advanced file options**.

Step 7 Click **Browse** to select the file that you saved in Step 1.

Step 8 From the Web Appliances drop-down list, choose **Select appliances in list** or **All assigned to Master** and then select the appliances to which you want to publish the configuration file.

Step 9 Click **Publish.**

Advanced File Publish: Publish Later

Procedure

Step 1 From the source Web Security appliance, save a Configuration File.

For configuration file compatibility information, see the SMA Compatibility Matrix, page 2-2.

For instructions on saving a configuration file from a Web Security appliance, see the *Cisco IronPort AsyncOS for Web Security User Guide*.

Step 2 On each destination Web Security appliance, save the existing configuration on your Web Security appliance to a configuration file. See the *Cisco IronPort AsyncOS for Web Security User Guide* for details.

Step 3 On the main Security Management appliance window, choose **Web > Utilities > Publish to Web Appliances**.

Step 4 Click **Schedule a Job.**

Step 5 **System-generated job name** is selected by default, or enter a job name (up to 80 characters).

Step 6 Enter the date and time when you want to publish the configuration.

Step 7 For **Configuration Master to Publish**, select **Advanced file options**, then click **Browse** to select the configuration file that you saved in Step 1.

Step 8 From the Web Appliances drop-down list, choose **Select appliances in list** or **All assigned to Master** and then select the appliances to which you want to publish the configuration file.

Step 9 Click **Publish.**

Viewing Status and History of Publishing Jobs

- Viewing Scheduled Publishing Jobs, page 8-18
- Viewing Status of the Current Publishing Job, page 8-18
- Viewing Publish History, page 8-18

Viewing Scheduled Publishing Jobs

To view a list of publishing jobs that are scheduled but have not yet occurred, choose **Web > Utilities > Publish to Web Appliances** and look in the **Pending Jobs** section.

Viewing Status of the Current Publishing Job

To view the status of a publishing job that is currently in progress, choose **Web > Utilities > Publish to Web Appliances** and look in the **Publishing Progress** section.

Viewing Publish History

Viewing the publish history is useful for checking for errors that may have occurred during publishing.

Procedure

Step 1 On the Security Management appliance, choose **Web > Utilities > Publish History**.

The Publish History page lists all of your most recent publish jobs attempted. Column information includes: Job Name, Job Completion Time, Configuration Master used (or name of the XML configuration file, if you performed an advanced file publish), Number of Appliances to which you published the job, and Status (Success or Failure).

Step 2 To view additional details about a particular job, click the specific job name hypertext link in the Job Name column.

Step 3 On the Publish History: Job Details page you can view additional details about a particular appliance in the job by clicking the appliance name; the **Web > Utilities > Web Appliance Status** page appears. You can also view status details about a particular appliance in the job, click the corresponding Details link to view the details on the Web Appliance Publish Details page.

Publishing fails if there is a discrepancy between the status in the “Web Appliance Service” column and the status in the “Is Service Displayed on Management Appliance?” column on the Web Appliance Status page. If both columns show that the feature is enabled but the corresponding Feature Key is not active (for example, is expired), publishing will also fail.

Viewing Web Security Appliance Status

Web Appliance Status Page

The **Web > Utilities > Web Appliance Status** page provides a high-level summary of the Web Security appliances connected to your Security Management appliance.

Note Only machines with support for centralized management will have data available for display.

Figure 8-1 Web Appliance Status Page

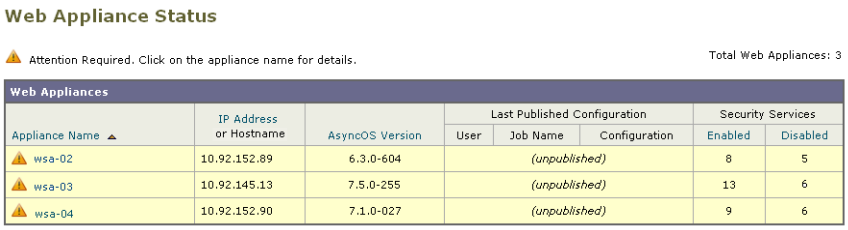

Note Warning messages will appear if different versions of the Acceptable Use Control Engine on the Web Security appliance do not match with those on the Security Management appliance. An 'N/A' is displayed if the service is disabled or not present on the Web Security appliance.

The Web Appliance Status page displays a list of your connected Web Security appliances, including appliance name, IP address, AsyncOS version, last published configuration information (user, job name, and configuration version), number of security services enabled or disabled, and total number of connected appliances (up to 150). The warning icon indicates when attention is required for one of your connected appliances.

Note It can take several minutes for the Web Appliance Status page to reflect recent configuration changes that occurred on the Web Security appliances. To refresh the data immediately, click the **Refresh Data** link. The time stamp on the page tells you when the data was last refreshed.

Appliance Status Page

The Appliance Status page provides a detailed view into the status of each connected appliance.

To view details for a managed Web Security appliance on the Web Appliance Status page, click the name of the appliance.

Status information includes general information about the connected Web Security appliances, their published configuration, publish history, feature key status, and so forth.

The Appliance Status page has several sections:

- Web Appliance Status Details page: System Status, Configuration Publishing History, and Centralized Reporting Status

- Web Appliance Status Details Page: Security Services Section
- Web Appliance Status Details Page: AnyConnect Secure Mobility, Proxy, and Authentication Settings

Figure 8-2 Web Appliance Status Details page: System Status, Configuration Publishing History, and Centralized Reporting Status

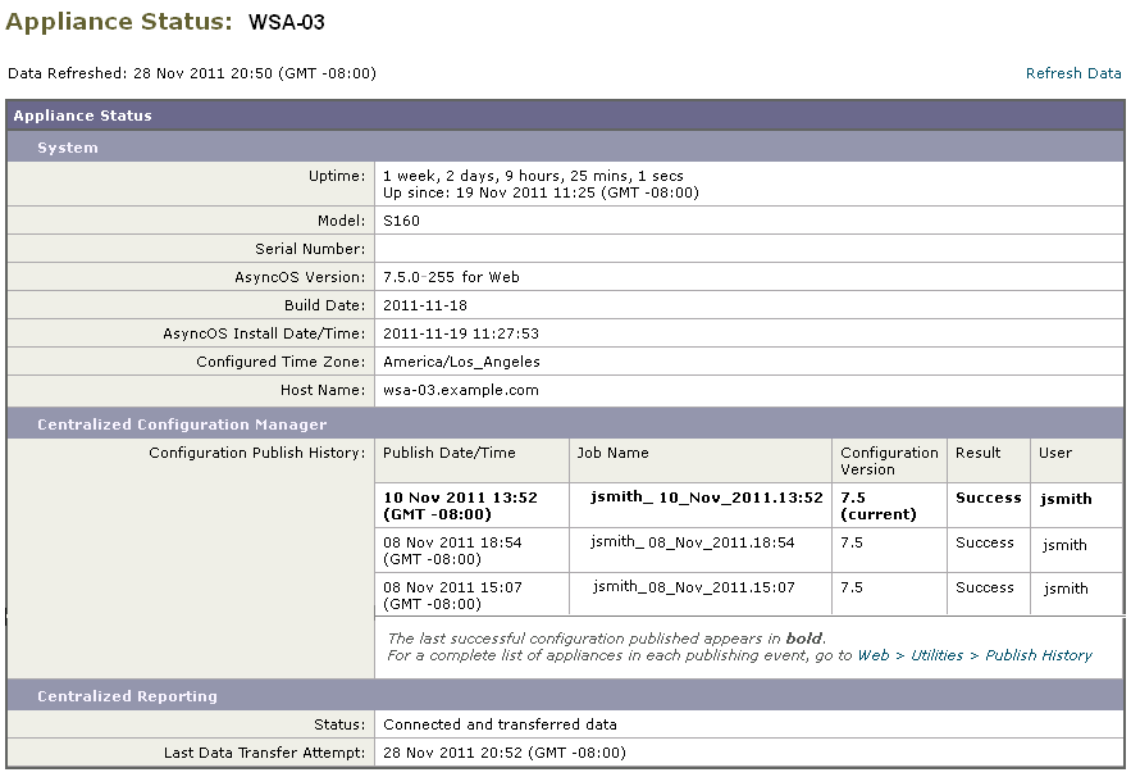

Appliance Status: WSA-03

Data Refreshed: 28 Nov 2011 20:50 (GMT -08:00) Refresh Data

Appliance Status

System

Uptime:	1 week, 2 days, 9 hours, 25 mins, 1 secs Up since: 19 Nov 2011 11:25 (GMT -08:00)
Model:	S160
Serial Number:	
AsyncOS Version:	7.5.0-255 for Web
Build Date:	2011-11-18
AsyncOS Install Date/Time:	2011-11-19 11:27:53
Configured Time Zone:	America/Los_Angeles
Host Name:	wsa-03.example.com

Centralized Configuration Manager

Configuration Publish History:

Publish Date/Time	Job Name	Configuration Version	Result	User
10 Nov 2011 13:52 (GMT -08:00)	**jsmith_ 10_Nov_2011.13:52**	**7.5 (current)**	**Success**	**jsmith**
08 Nov 2011 18:54 (GMT -08:00)	jsmith_ 08_Nov_2011.18:54	7.5	Success	jsmith
08 Nov 2011 15:07 (GMT -08:00)	jsmith_08_Nov_2011.15:07	7.5	Success	jsmith

*The last successful configuration published appears in **bold**.*
For a complete list of appliances in each publishing event, go to Web > Utilities > Publish History

Centralized Reporting

Status:	Connected and transferred data
Last Data Transfer Attempt:	28 Nov 2011 20:52 (GMT -08:00)

Figure 8-3 Web Appliance Status Details Page: Security Services Section

Security Services

⚠ *One or more of the services on the Web Appliance does not match the corresponding Security Service Display setting on the Management Appliance.*

Description	Services		Feature Keys		
	Web Appliance Service	Is Service Displayed on Management Appliance?	Status	Time Remaining	Expiration Date
Cisco IronPort Web Proxy & DVS(TM) Engine	Enabled	Yes	Active	Perpetual	N/A
Cisco IronPort L4 Traffic Monitor	N/A	N/A	N/A	N/A	N/A
Proxy Mode	Forward	No (Bypass Proxy)			
FTP Proxy	Enabled	Yes			
Cisco IronPort HTTPS Proxy	Enabled	Yes	Active	Perpetual	N/A
SOCKS Proxy	Disabled	Yes			
Upstream Proxy Groups	Not Configured	No (Routing Policies)			
AnyConnect Secure Mobility	Cisco ASA	Yes (Cisco ASA)	Active	Perpetual	N/A
Cisco IronPort Web Usage Controls	Enabled	Yes	Active	Perpetual	N/A
Application Visibility and Control	Enabled	Yes			
Cisco IronPort Centralized Web Reporting	Enabled	Yes			
Cisco IronPort Web Reputation Filters	Enabled	Yes	Active	Perpetual	N/A
Adaptive Scanning	Disabled	No			
Webroot Anti-Malware	Enabled	Yes	Active	Perpetual	N/A
McAfee Anti-Malware	Enabled	Yes	Active	Perpetual	N/A
Sophos Antivirus	Enabled	Yes	Active	Perpetual	N/A
End-User Acknowledgement	Enabled	Yes			
Cisco IronPort Data Security Filters	Enabled	Yes			
External DLP Servers	Not Configured	No			
Credential Encryption	Disabled	No			
Identity Provider for SaaS	Not Configured	No			

Acceptable Use Controls Engine Updates

Update Type	Web Appliance Version	Management Appliance Version
Web Categorization Categories List	1337225318	1337225318
Application Visibility and Control Data	1346296993	1346296993

Figure 8-4 Web Appliance Status Details Page: AnyConnect Secure Mobility, Proxy, and Authentication Settings

AnyConnect Secure Mobility Settings

Cisco ASA:	[IP address and port go here]

Proxy Settings

Upstream Proxies:	*No upstream proxies configured.*
HTTP Ports to Proxy:	[Port goes here]

Authentication Service

Authentication Realms:	Name	Protocol	Servers	Support Transparent User Identification
	NTLM AUTH	NTLM	ad.example.com	No
	LDAP AUTH	LDAP	ad.example.com	No
Authentication Sequences:	Name	Order of Realms		
	All Realms	NTLMSSP: NTLM AUTH Basic: NTLM AUTH, LDAP AUTH		
Unreachable Authentication Service Action:	Block all traffic if authentication fails			

Details include:

- System status information (uptime, appliance model and serial number, AsyncOS version, build date, AsyncOS installation date and time, and host name)
- Configuration publish history (publish date/time, job name, configuration version, result of the publish, and user)
- Centralized reporting status, including time of last attempted data transfer
- Web Security feature status (feature description, configuration summary, security services settings, and status of feature keys)
- Use this information to when configuring your appliances for Centralized Management.
- Acceptable Use Controls Engine versions on the managed and managing appliances
- AnyConnect Secure Mobility settings
- Proxy settings (upstream proxies and HTTP ports to proxy)
- Authentication service (name, protocol, and servers of authentication realms; name and order of realms in authentication sequences; whether Transparent User Identification is supported; and whether to block or permit traffic if authentication fails)

Tip To refresh the details, if for example you have made changes on the Web Security appliance, or if a message appears indicating that information for the appliance is not yet available, click the **Refresh Data** link. The time stamp on the page tells you when the data was last refreshed.

URL Category Set Updates and Centralized Configuration Management

URL category set updates apply only to Configuration Master 7.5 and higher.

In order to ensure that your system has the latest set of predefined URL categories available for managing web usage, the URL category set for Cisco IronPort Web Usage Controls (WUC) may be updated occasionally: By default, Web Security appliances download URL category set updates automatically from Cisco IronPort, and the Security Management appliance receives these updates automatically within a few minutes from managed Web Security appliances. The updated set of URL Categories will appear immediately in identities and policies in Configuration Master 7.5 and higher.

Actions you should take include the following:

- Understand the Impacts of URL Category Set Updates, page 8-22
- Ensure that You Will Receive Alerts about URL Category Set Updates, page 8-23
- Be Aware: Before You Set Up Configuration Master 7.5 and 7.7, page 8-23
- Specify Default Settings for New and Changed Categories, page 8-23
- When the URL Category Set is Updated, Check Your Policy and Identity Settings, page 8-23

Understand the Impacts of URL Category Set Updates

For essential information about actions you should take before and after URL category set updates, see the “Managing Updates to the Set of URL Categories” section of the “URL Filters” chapter in the *Cisco IronPort AsyncOS for Web Security User Guide* at the link provided in Documentation, page 1-2. Category descriptions are in the “URL Category Descriptions” section of the same chapter.

Ensure that You Will Receive Alerts about URL Category Set Updates

To ensure that you will receive alerts about URL category set updates that will affect policy settings in your Configuration Master, go to Management Appliance > System Administration > Alerts and make sure you will receive Warning-level alerts in the System category. More information about alerts is in Managing Alerts, page 13-28.

Be Aware: Before You Set Up Configuration Master 7.5 or 7.7

If you copy or import Configuration Master 7.1 settings into Configuration Master 7.5 or 7.7, identities and policies that reference URL categories will be modified in those Configuration Masters. You should import a configuration file from a properly-configured Web Security appliance instead, or evaluate the Uncategorized URLs settings in each policy in Configuration Master 7.1 before copying or importing.

Specify Default Settings for New and Changed Categories

If URL category set updates occur in future, they may change the behavior of existing policies. Before URL category set updates occur, you should specify default actions for new and merged categories in each policy that offers URL filtering, or import a configuration from a Web Security appliance that has these settings already configured.

For more information, see the "Choosing Default Settings for New and Changed Categories" section in the "URL Filters" chapter of the *User Guide for Cisco IronPort AsyncOS for Web Security* or the online help on the Web Security appliance.

When the URL Category Set is Updated, Check Your Policy and Identity Settings

Category set updates trigger two types of alerts:

- Alerts about category changes
- Alerts about policies that have changed or been disabled as a result of category changes

When you receive alerts about URL category set changes, you should check your existing URL category-based policies and identities in Configuration Master 7.5 and higher to be sure they still meet your policy goals.

For more information about the kinds of changes that might require your attention, see the "Responding to Alerts about URL Category Set Updates" section in the *Cisco IronPort AsyncOS for Web Security User Guide* at the link provided in Documentation, page 1-2.

CHAPTER 9

Monitoring System Status

- Monitoring Security Management Appliance Status, page 9-1
- Viewing the Status of Your Managed Appliances, page 9-6
- Monitoring Capacity of Managed Appliances, page 9-7
- Monitoring Reporting Data Availability Status, page 9-7
- Monitoring Tracking Data Status, page 9-8
- Identifying Active TCP/IP Services, page 9-10

Monitoring Security Management Appliance Status

By default, the System Status page is the first page that appears when you access the GUI of the Security Management appliance. (To change the landing page, see Setting Preferences, page 13-52.)

To access the System Status page from anywhere in the GUI, select **Management Appliance > Centralized Services > System Status**.

Figure 9-1 System Status Page

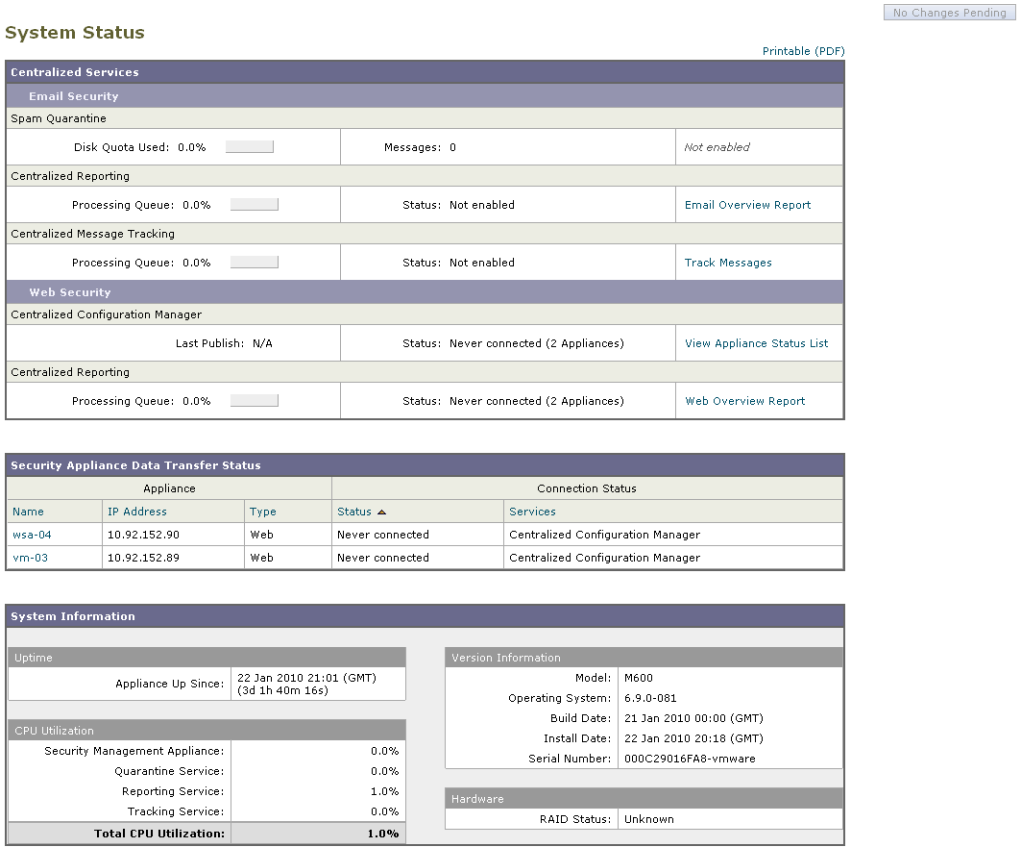

Before you enable monitoring services and add a managed appliance, only the System Information section provides status information. If you have run the System Setup Wizard, enabled monitoring services, and added a managed appliance, the Services section and the Security Appliance Data Transfer Status section are populated with data. For information about adding managed appliances, enabling services, and viewing the status of both, see About Adding Managed Appliances, page 2-12, Configuring Services on the Security Management Appliance, page 2-14, and Viewing the Status of Your Managed Appliances, page 9-6.

In general, the System Status page provides detailed status information for the Security Management appliance, including the following:

- **System Status:** overview of the services (Cisco IronPort Spam Quarantine, centralized reporting, centralized tracking, and centralized configuration manager)
- **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 version, build date, installation date, and serial number

Centralized Services

The Centralized Services section provides an overview of services on the Security Management appliance.

On the main Security Management appliance, choose **Management Appliance > Centralized Services > System Status**, to view summary information about the transfer of reporting data between the managed Email Security appliances, Web Security appliances, and the Security Management appliance.

There are two sections under **Centralized Services > System Status**:

- Email Security
- Web Security

Email Security

The Email Security section pertains to information for the Email Security appliances only. The Email Security section displays information for:

- **Spam Quarantine**—In this section, you can see the number of messages held in the Cisco IronPort Spam Quarantine and the percentage of the disk quota used by the quarantine. Click the Spam Quarantine View link to access the Spam Quarantine page. For more information about the Cisco IronPort Spam Quarantine, see Chapter 7, "Managing the Cisco IronPort Spam Quarantine.".
- **Centralized Reporting**—In this section you can see the Processing Queue information which shows the percentage of the processing queue that is used by reporting data.

  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 Email Security appliances and processed by the Security Management appliance. If the processing queue usage percentage remains 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.

Note The 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.

  Click the Email Overview Report link to access the Overview interactive report page. For more information about the Overview report page, see Email Reporting Overview Page, in Chapter 4, "Using Centralized Email Security Reporting."

- **Centralized Message Tracking**—In this section you can view summary information about the transfer of tracking data between the managed Email Security appliances and the Security Management appliance. The Processing Queue field shows what percentage of the processing queue is taken up by tracking data. Click the Track Messages link to access the Message Tracking query page. For more information about tracking messages, see Chapter 6, "Tracking Email Messages."

Web Security

The Web Security section pertains to information for the Web Security appliances only. The Web Security section displays information for:

- **Centralized Reporting**—In this section you can see the Processing Queue information which shows the percentage of the processing queue that is used by reporting data.

  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 Web Security appliances and processed by the Security Management appliance. If the processing queue usage percentage remains 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.

Note The 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.

Click the Web Overview Report link to access the Overview interactive report page. For more information about the Overview report page, see Understanding the Web Reporting Pages, in Chapter 5, "Using Centralized Web Reporting and Tracking."

- **Centralized Configuration Manager**—In this section you can view summary information about the last successful configuration update of a Web Security appliance. You an click the interactive link that takes you to the last successfully published update on your system. Click the View Appliance Status List to view the status of individual appliances on the Security Management appliance. For more information the status of appliances currently on your system, see Viewing the Status of Your Managed Appliances, page 9-6.

Security Appliance Data Transfer Status

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, you can view connection status issues for specific appliances. For detailed information, click an appliance name to view the Data Transfer Status page for the appliance.

Figure 9-2 Data Transfer Status: <Appliance_Name> Page

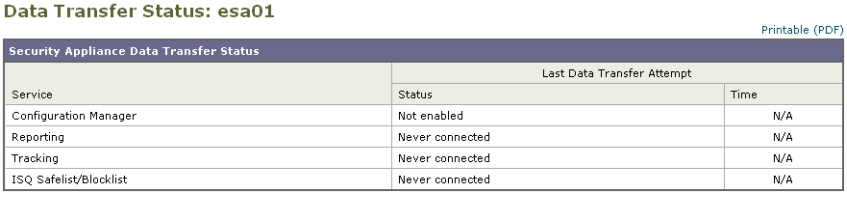
Data Transfer Status: esa01

Printable (PDF)

Security Appliance Data Transfer Status		
	Last Data Transfer Attempt	
Service	Status	Time
Configuration Manager	Not enabled	N/A
Reporting	Never connected	N/A
Tracking	Never connected	N/A
ISQ Safelist/Blocklist	Never connected	N/A

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.

The data transfer status for Web Security appliances can be one of the following values:

- **Not enabled:** The centralized configuration manager is not enabled for the Web Security appliance.
- **Never connected:** The centralized configuration manager is enabled for the Web Security appliance, but no connection has been established between the Web Security appliance and the Security Management appliance.
- **Waiting for data:** The Web 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 Web Security appliance and the Security Management appliance, and data were successfully transferred.
- **Configuration push failure:** The Security Management appliance attempted to push a configuration file to the Web Security appliance, but the transfer failed.
- **Configuration push pending:** The Security Management appliance is in the process of pushing a configuration file to the Web Security appliance.
- **Configuration push success:** The Security Management appliance successfully pushed a configuration file to the Web Security appliance.

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.

System Information

The System Information section of the System Status page provides information about the operating system and performance of the Security Management appliance. The Uptime field shows when the appliance was last started and how long it has been running. The Version Information area lists the model number, the version of AsyncOS, the build and install dates of the operating system, and the serial number of the appliance.

Note You may need to supply the appliance serial number to assist Cisco IronPort Customer Support with troubleshooting.

Figure 9-3 System Information Section of the System Status Page

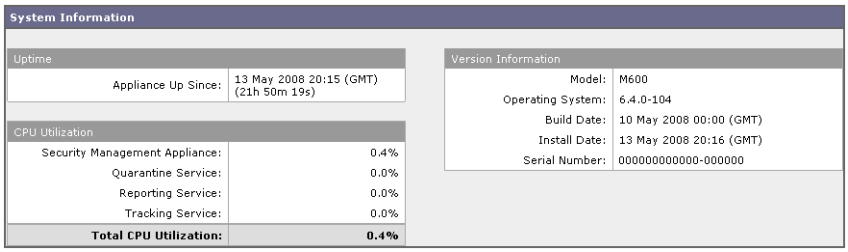

The CPU Utilization percentages indicate the portion of the Security Management appliance's CPU processing that is devoted to each of the monitoring services. The percentages show the amount of CPU currently used by the three major services. Other operations of the Security Management appliance are grouped under the general heading "Security Management appliance."

The CPU utilization percentages change constantly. Refresh the browser display to view the most recent data.

Viewing the Status of Your Managed Appliances

The Security Appliances page provides information about the status of your managed appliances.

To access the Security Appliances page, perform the following:

On the Security Management appliance, choose **Management Appliance > Centralized Services > Security Appliances**.

Figure 9-4 Security Appliances Page

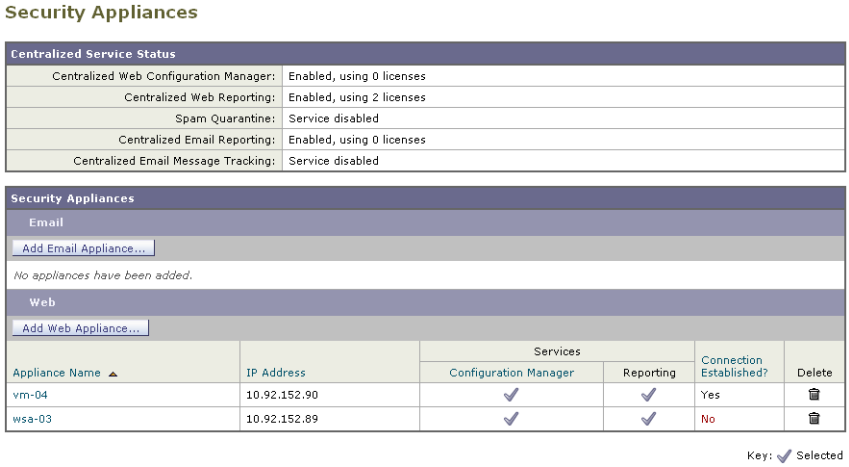

The Centralized Service Status section shows which services are enabled and how many licenses you have used for each service. The Security Appliances section lists the appliances you have added. Check marks indicate the enabled services, and the Connection Established? column shows whether or not file transfer access is properly configured. You can also add or delete an appliance. For more information, see About Adding Managed Appliances, page 2-12.

Additional Status Information for Web Security Appliances

For additional status information about Web Security appliances, see Viewing Web Security Appliance Status, page 8-19.

Monitoring Capacity of Managed Appliances

To View Capacity of	See
Web Security appliances	System Capacity Page, page 5-55
Email Security appliances	System Capacity Page, page 4-43

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 Data Availability on your Email Security Appliances, page 9-7
- Monitoring Data Availability on your Web Security Appliances, page 9-8

Monitoring Data Availability on your Email Security Appliances

To monitor reporting data from your Email Security appliances on the Security Management appliance, view the **Email > Reporting > Reporting Data Availability** page.

Figure 9-5 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.

Note If you have had to replace an Email Security appliance due to a hardware failure or other reasons, the data from the replaced Email Security appliance will not be lost, but the data will not be displayed correctly on the Security Management appliance.

Monitoring Data Availability on your Web Security Appliances

To monitor reporting data from your Web Security appliances on the Security Management appliance, view the **Web > Reporting > Data Availability** page.

From the Data Availability page you can update and sort data to provide real-time visibility into resource utilization and web traffic trouble spots.

Web Reporting Data Availability

Printable (PDF)

Web Reporting Data Range

Displaying 1 - 2 of 2 appliances.

Web Security Appliance	Web Reporting		Web Tracking and Reporting Detail		Status
	From	To	From	To	
vmw098-wsa08.sma	26 Aug 2010 09:00	27 Aug 2010 02:22	26 Aug 2010 11:00	27 Aug 2010 02:22	Ok
vmw095-wsa11.sma	N/A	N/A	N/A	N/A	Never Connected
Overall:	26 Aug 2010 09:00 (GMT +03:00)	27 Aug 2010 02:22 (GMT +03:00)	26 Aug 2010 11:00 (GMT +03:00)	27 Aug 2010 02:22 (GMT +03:00)	

Displaying 1 - 2 of 2 appliances.

Note In the Web Reporting Data Availability window, Web Reporting will show disabled only if **both** Web Reporting and Email Reporting are disabled.

All data resource utilization and web traffic trouble spots are shown from this page. By clicking on one of the listed Web Security appliance links, you can view reporting data availability for that appliance.

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 Web 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.

If data availability is used within a scheduled report for URL Categories, and there are gaps in data for any of the appliances, the following message is displayed at the bottom of the page: “Some data in this time range was unavailable.” If there are no gaps present, nothing appears.

See the “Data Availability Page” for more information on the Data Availability page on the Web Security appliance.

Monitoring Tracking Data Status

The Security Management appliance allows you to monitor and track data from any appliance.

- Monitoring Email Tracking Data Status, page 9-9
- Monitoring Web Tracking Data Status, page 9-9

Monitoring Email Tracking Data Status

To monitor the status of email tracking data, view the **Email > Message Tracking > Message Tracking Data Availability** page.

Note The Email Security appliance makes duplicate copies of reporting and tracking data taken from that appliance and places copies of the data files into additional folders apart from the default directory. The Security Management appliance can then be configured to pull data from one of those folders.

Figure 9-6 Message Tracking Data Availability Page

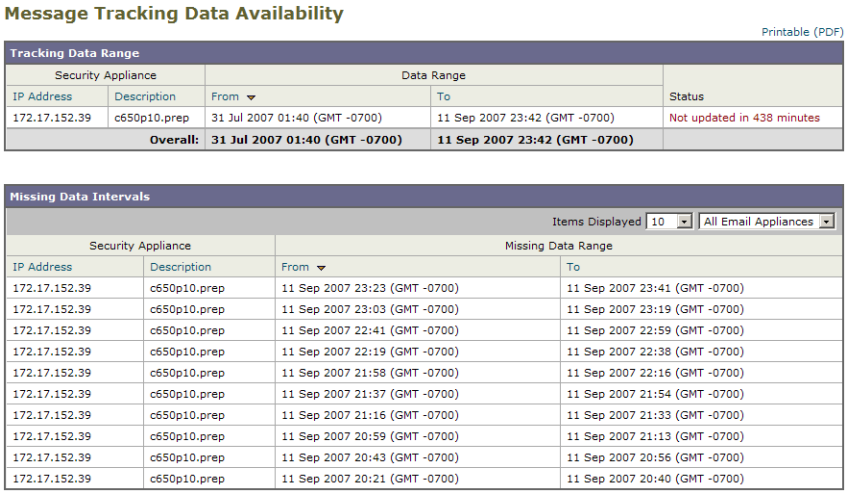

Message Tracking Data Availability

Printable (PDF)

Tracking Data Range

Security Appliance		Data Range		
IP Address	Description	From	To	Status
172.17.152.39	c650p10.prep	31 Jul 2007 01:40 (GT -0700)	11 Sep 2007 23:42 (GT -0700)	Not updated in 438 minutes
	Overall:	31 Jul 2007 01:40 (GT -0700)	11 Sep 2007 23:42 (GT -0700)	

Missing Data Intervals

Items Displayed 10 All Email Appliances

Security Appliance		Missing Data Range	
IP Address	Description	From	To
172.17.152.39	c650p10.prep	11 Sep 2007 23:23 (GT -0700)	11 Sep 2007 23:41 (GT -0700)
172.17.152.39	c650p10.prep	11 Sep 2007 23:03 (GT -0700)	11 Sep 2007 23:19 (GT -0700)
172.17.152.39	c650p10.prep	11 Sep 2007 22:41 (GT -0700)	11 Sep 2007 22:59 (GT -0700)
172.17.152.39	c650p10.prep	11 Sep 2007 22:19 (GT -0700)	11 Sep 2007 22:38 (GT -0700)
172.17.152.39	c650p10.prep	11 Sep 2007 21:58 (GT -0700)	11 Sep 2007 22:16 (GT -0700)
172.17.152.39	c650p10.prep	11 Sep 2007 21:37 (GT -0700)	11 Sep 2007 21:54 (GT -0700)
172.17.152.39	c650p10.prep	11 Sep 2007 21:16 (GT -0700)	11 Sep 2007 21:33 (GT -0700)
172.17.152.39	c650p10.prep	11 Sep 2007 20:59 (GT -0700)	11 Sep 2007 21:13 (GT -0700)
172.17.152.39	c650p10.prep	11 Sep 2007 20:43 (GT -0700)	11 Sep 2007 20:56 (GT -0700)
172.17.152.39	c650p10.prep	11 Sep 2007 20:21 (GT -0700)	11 Sep 2007 20:40 (GT -0700)

The Message Tracking Data Availability page allows you to view the missing-data intervals for the Security Management appliance. A missing-data interval is a period of time during which the Security Management appliance received no message tracking data from your organization's Email Security appliances.

You can monitor data availability for a particular managed appliance or for all Email Security appliances in the system. If you find missing-data intervals in the message tracking data, you can immediately tell that your data may be incomplete. Use the data availability information to validate your message tracking data and to troubleshoot system problems.

Monitoring Web Tracking Data Status

Procedure

Step 1 On the Security Management appliance, choose **Web > Reporting > Web Tracking**.

Step 2 From the **Time Range** drop-down list, select the time range that you want information about.

Step 3 In the User/Client IP or the Website text field, enter a value.

Step 4 From the Transaction Type drop-down list, select the type of transaction.

Choices include: All Transactions, Completed, Blocked, Monitored, Warned.

The following example shows the results when 'google.com' is entered in the Website text field.

Figure 9-7 Web Tracking Data Status Page

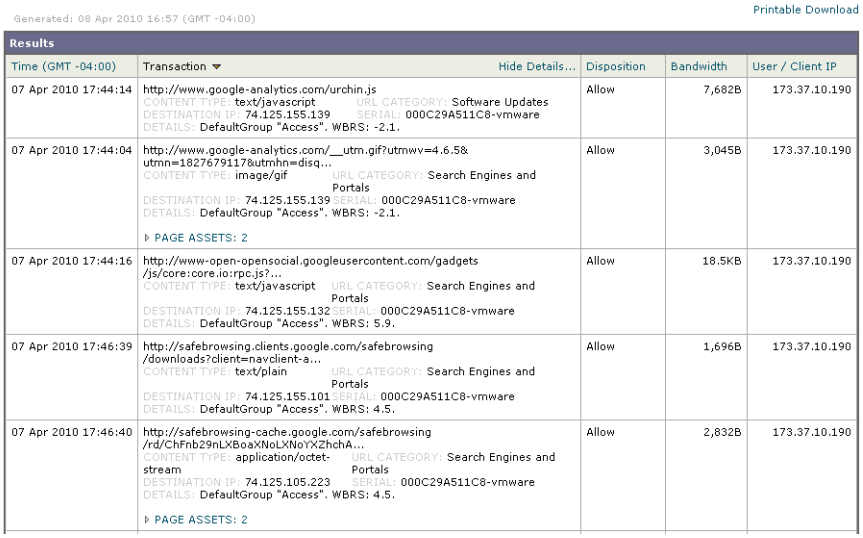

For more information on Web Tracking, see the Web Tracking, page 5-50.

Identifying Active TCP/IP Services

To identify active TCP/IP services, use the `tcpservices` command in the command line interface.

CHAPTER 10

Integrating with LDAP

- Overview, page 10-1
- Configuring LDAP to Work with the Cisco IronPort Spam Quarantine, page 10-1
- Creating the LDAP Server Profile, page 10-2
- Configuring LDAP Queries, page 10-4
- Domain-Based Queries, page 10-8
- Chain Queries, page 10-10
- Configuring AsyncOS to Work With Multiple LDAP Servers, page 10-11
- Configuring External Authentication of Administrative Users Using LDAP, page 10-14

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 Cisco IronPort Spam Quarantine.

  When a user logs in to the web UI for the Cisco IronPort 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 Cisco IronPort Spam Quarantine, as long as the appliance does not rewrite them.

  See Configuring LDAP to Work with the Cisco IronPort Spam Quarantine, page 10-1.

- Administrative users who sign in to the Security Management appliance when External Authentication is enabled and configured.

  See Configuring External Authentication of Administrative Users Using LDAP, page 10-14.

Configuring LDAP to Work with the Cisco IronPort Spam Quarantine

When you configure your Cisco IronPort appliance to work with an LDAP directory, you must complete the following steps to set up for acceptance, routing, aliasing, and masquerading:

Procedure

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, page 10-2.

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, page 10-11.

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, page 10-4.

Step 3 **Enable LDAP end-user access and spam notifications for the Cisco IronPort Spam Quarantine**.

Enable LDAP end-user access to the Cisco IronPort 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, page 7-2.

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.

Procedure

Step 1 On the main Security Management appliance, choose **Management Appliance > System Administration > LDAP**.

Step 2 Click **Add LDAP Server Profile.**

Step 3 Enter a name for the server profile in the **LDAP Server Profile Name** text field.

Step 4 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, page 10-11.

Step 5 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 6 Select the LDAP server type: Active Directory, OpenLDAP, or Unknown or Other.

Step 7 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 8 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 9 Under Advanced, select whether to use SSL when communicating with the LDAP server.

Step 10 Enter the cache time-to-live. This value represents the amount of time to retain caches.

Step 11 Enter the maximum number of retained cache entries.

Step 12 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, page 10-13.

Note The maximum number of simultaneous connections includes LDAP connections used for LDAP queries. However, if you enable LDAP authentication for the Cisco IronPort Spam Quarantine, the appliance allows 20 additional connections for the end user quarantine for a total of 30 connections.

Step 13 Test the connection to the server by clicking the **Test Server(s)** button. If you specified multiple LDAP servers, they are all tested. The results of the test appear in the Connection Status field. For more information, see Testing LDAP Servers, page 10-4.

Step 14 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, page 10-4.

Step 15 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 16 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 Cisco IronPort Spam Quarantine query:

- **Spam quarantine end-user authentication query.** For more information, see the "Spam Quarantine End-User Authentication Queries" section on page 10-5.
- **Spam quarantine alias consolidation query.** For more information, see Spam Quarantine Alias Consolidation Queries, page 10-6.

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, page 10-8 and Chain Queries, page 10-10.

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" section on page 10-8 for more information.

Spam Quarantine End-User Authentication Queries

End-user authentication queries validate users when they log in to the Cisco IronPort 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 10-1 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 10-2 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 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 IronPort 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 10-3 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 10-4 Example LDAP Server and Spam Quarantine Alias Consolidation Settings: OpenLDAP

Authentication Method	Anonymous
Server Type	OpenLDAP
Port	389
Base DN	[Blank] (Some older schemas will want to use a specific Base DN.)
Connection Protocol	Use SSL
Query String	`(mail={a}))`
Email Attribute	`mail`

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 IronPort 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 Cisco IronPort Spam Quarantine, complete the following steps:

Procedure

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, page 10-2.

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 Cisco IronPort Spam Quarantine. For more information about creating the query, see Creating a Domain-Based Query, page 10-9.

Step 3 Enable end-user access or spam notifications for the Cisco IronPort Spam Quarantine. For more information, see Setting Up the Centralized Spam Quarantine, page 7-2.

Creating a Domain-Based Query

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > System Administration > LDAP**.

Step 2 On the LDAP page, click **Advanced**.

Step 3 Enter a name for the domain-based query.

Step 4 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 5 In the Domain Assignments field, enter a domain.

Step 6 Select a query to associate with the domain.

Step 7 Add a row and select a query for each domain in the domain-based query.

Step 8 Enter a default query to run if all other queries fail. If you do not want to enter a default query, select **None**.

Figure 10-1 Example Domain-Based Query

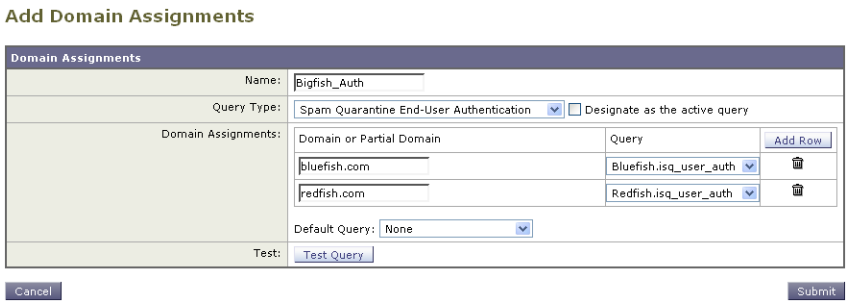

Step 9 Test the query by clicking the **Test Query** button and entering a user login and password or an email address to test in the Test Parameters fields. The results appear in the Connection Status field.

Step 10 Check the **Designate as the active query** checkbox if you want the Cisco IronPort 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 Cisco IronPort Spam Quarantine.

Step 11 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 Cisco IronPort Spam Quarantine, complete the following steps:

Procedure

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, page 10-2.

Step 2 Create the chain query and designate it as an active query for the Cisco IronPort Spam Quarantine. For more information, see Creating a Chain Query, page 10-10.

Step 3 Enable LDAP end-user access or spam notifications for the Cisco IronPort Spam Quarantine. For more information about the spam quarantine, see Setting Up the Centralized Spam Quarantine, page 7-2.

Creating a Chain Query

Tip You can also use the `advanced` subcommand of the `ldapconfig` command in the CLI.

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > System Administration > LDAP > LDAP Server**.

Step 2 From the LDAP Server Profiles page, click **Advanced**.

Step 3 Click **Add Chained Query**.

Step 4 Enter a name for the chain query.

Step 5 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 6 Select the first query in the chain.

The Cisco IronPort 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 10-2 Example Chain Query

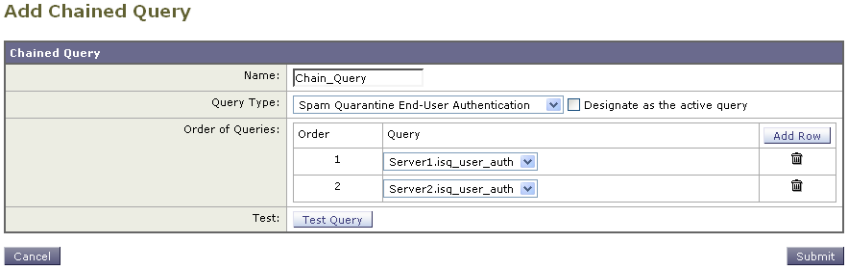

Step 7 Test the query by clicking the **Test Query** button and entering a user login and password or an email address in the Test Parameters fields. The results appear in the Connection Status field.

Step 8 Check the **Designate as the active query** check box if you want the Cisco IronPort 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 Cisco IronPort Spam Quarantine.

Step 9 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 IronPort 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 IronPort appliance to connect to redundant LDAP servers to use the following features:

- **Failover**. If the Cisco IronPort appliance cannot connect to an LDAP server, it connects to the next server in the list.
- **Load Balancing**. The Cisco IronPort 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.

The Cisco IronPort appliance attempts to connect to the first server in the list of LDAP servers for a specified period of time. If the appliance cannot connect to the first LDAP server in the list, the appliance attempts to connect to the next LDAP server in the list. To ensure that the Cisco IronPort appliance connects to the primary LDAP server by default, enter it as the first server in the list of LDAP servers.

If the Cisco IronPort 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 IronPort Appliance for LDAP Failover

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > System Administration > LDAP**.

Step 2 Select the LDAP server profile you want to edit.

In the following example, the LDAP server name is example.com.

Figure 10-3 Example LDAP Failover Configuration

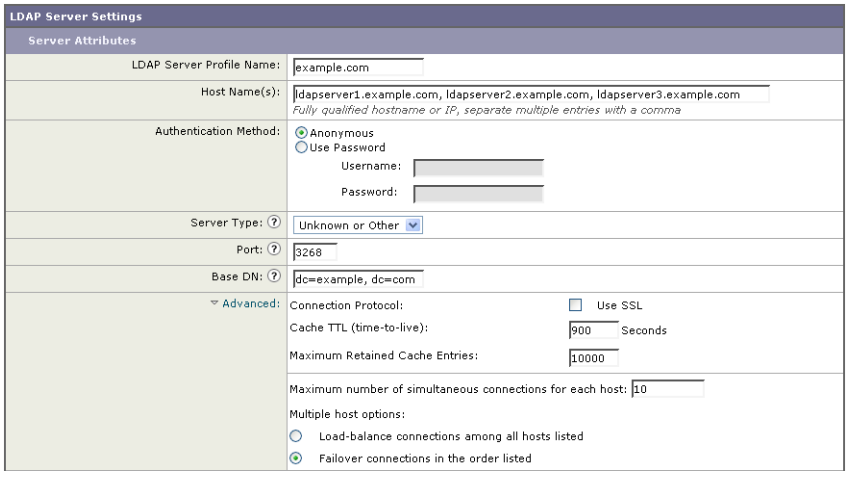

Step 3 In the Hostname text field, type the LDAP Servers; for example **ldapserver.example.com**.

Step 4 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 5 Click on the radio button next to **Failover connections in the order list**.

Step 6 Configure other LDAP options as necessary.

Step 7 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 IronPort appliance distributes connections among the LDAP servers listed. If a connection fails or times out, the Cisco IronPort appliance determines which LDAP servers are available and reconnects to available servers. The Cisco IronPort 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 Cisco IronPort appliance distributes the connection load among the remaining LDAP servers.

Configuring the Cisco IronPort Appliance for Load Balancing

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > System Administration > LDAP**.

Step 2 Select the LDAP server profile you want to edit

In the following example, the LDAP server name is example.com.

Figure 10-4 Example Load Balancing Configuration

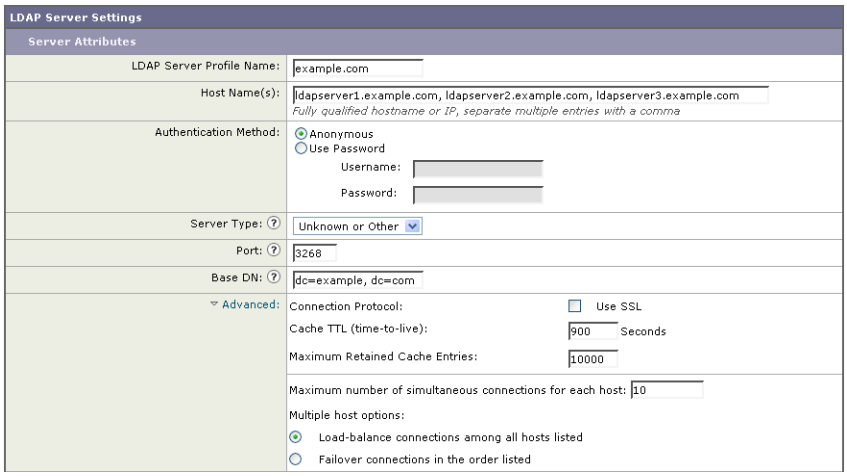

Step 3 In the Hostname text field, type the LDAP Servers; for example **ldapserver.example.com**.

Step 4 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 5 Click on the radio button next to **Load balance connections among all hosts**.

Step 6 Configure other LDAP options as necessary.

Step 7 Submit and commit the changes.

Configuring External Authentication of Administrative Users Using LDAP

You can configure the Cisco IronPort appliance to use an LDAP directory on your network to authenticate administrative users by allowing them to log in to the Security Management appliance with their LDAP user names and passwords.

Procedure

Step 1 **Configure the LDAP Server Profile.** See Creating the LDAP Server Profile, page 10-2.

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, page 10-15.

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, page 10-15 and the *Cisco IronPort AsyncOS for Email Security Advanced User Guide*.

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, page 10-8.

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, page 10-17 and the "Adding Users" in the *Cisco IronPort AsyncOS for Email Security Advanced User Guide*.

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.

Table 10-5 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 10-5 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`

Table 10-6 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 10-6 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}.

Table 10-7 shows the default query strings and attributes that AsyncOS uses when it searches for group membership information on an Active Directory server.

Table 10-7 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`

Table 10-8 shows the default query strings and attributes that AsyncOS uses when it searches for group membership information on an OpenLDAP server.

Table 10-8 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:

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > System Administration > Users** page.

Step 2 Click **Enable**.

Step 3 Select the **Enable External Authentication** check box.

Step 4 Select **LDAP** for the authentication type.

Step 5 Select the LDAP external authentication query that authenticates users.

Step 6 Enter the number of seconds that the appliance waits for a response from the server before timing out.

Step 7 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 8 Optionally, click **Add Row** to add another directory group. Repeat steps 7 and 8 for each directory group that the appliance authenticates.

Step 9 Submit and commit your changes.

CHAPTER 11

Configuring SMTP Routing

This chapter explains the features that affect routing and delivery of email traveling through the Security Management appliance, and use of the SMTP Routes page and **smtproutes** command.

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.

SMTP Routes Overview

SMTP Routes allow you to redirect all email for a particular domain to a different mail exchange (MX) host. For example, you could make a mapping from `example.com` to `groupware.example.com`. This mapping causes any email with `@example.com` in the Envelope Recipient address to go instead to `groupware.example.com`. The system performs an "MX" lookup on `groupware.example.com`, and then performs an "A" lookup on the host, just like a normal email delivery. This alternate MX host does not need to be listed in DNS MX records and it does not even need to be a member of the domain whose email is being redirected. The Cisco IronPort AsyncOS operating system allows up to ten thousand (10,000) SMTP Route mappings to be configured for your Cisco IronPort appliance. (See SMTP Routes Limits, page 11-3.)

This feature also allows host "globbing." If you specify a partial domain, such as `example.com`, then any domain ending in `example.com` matches the entry. For instance, `fred@foo.example.com` and `wilma@bar.example.com` both match the mapping.

If a host is not found in the SMTP Routes table, an MX lookup is performed using DNS. The result is not re-checked against the SMTP Routes table. If the DNS MX entry for `foo.domain` is `bar.domain`, any email sent to `foo.domain` is delivered to the host `bar.domain`. If you create a mapping for `bar.domain` to some other host, email addressed to `foo.domain` is not affected.

In other words, recursive entries are not followed. If there is an entry for `a.domain` to redirect to `b.domain`, and a subsequent entry to redirect email for `b.domain` to `a.domain`, a mail loop will *not* be created. In this case, email addressed to `a.domain` will be delivered to the MX host specified by `b.domain`, and conversely email addressed to `b.domain` will be delivered to the MX host specified by `a.domain`.

The SMTP Routes table is read from the top down for every email delivery. The most specific entry that matches a mapping wins. For example, if there are mappings for both `host1.example.com` and `example.com` in the SMTP Routes table, the entry for `host1.example.com` will be used because it is the more specific entry — even if it appears after the less specific `example.com` entry. Otherwise, the system performs a regular MX lookup on the domain of the Envelope Recipient.

Default SMTP Route

You can also define a default SMTP route with the special keyword `ALL`. If a domain does not match a previous mapping in the SMTP Routes list, it defaults to being redirected to the MX host specified by the `ALL` entry.

When you print the SMTP Routes entries, the default SMTP route is listed as `ALL:`. You cannot delete the default SMTP route; you may only clear any values entered for it.

Configure the default SMTP route using the **Management Appliance > Network > SMTP Routes** page or the **smtproutes** command.

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.

SMTP Routes and DNS

Use the special keyword USEDNS to tell the appliance to do MX lookups to determine next hops for specific domains. This is useful when you need to route mail for subdomains to a specific host. For example, if mail to example.com is to be sent to the company's Exchange server, you might have something similar to the following SMTP route:

```
example.com exchange.example.com
```

However, for mail to various subdomains (foo.example.com), add an SMTP route that looks like this:

```
.example.com USEDNS
```

SMTP Routes, 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 IronPort appliance sits behind a mail relay server that is at the edge of the network.

Managing SMTP Routes on the Security Management Appliance

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > Network > SMTP Routes**.

Use this page to manage SMTP Routes on your Cisco IronPort appliance. From this page you can add, modify, and delete mappings in the table. You can export or import the SMTP Routes entries.

Adding SMTP Routes

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > Network > SMTP Routes**.

Step 2 Click **Add Route**.

Step 3 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 4 You can specify a port number by adding ":*<port number>*" to the destination host: `example.com:25`

Step 5 Submit and commit your changes.

Editing SMTP Routes

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > Network > SMTP Routes**.

Step 2 Click the name of an existing SMTP Route in the SMTP Route listing.

Step 3 Edit the route.

Step 4 Submit and commit your changes.

Deleting SMTP Routes

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > Network > SMTP Routes**.

Step 2 Select the check boxes to the right of the SMTP Routes to delete.

Step 3 Click **Delete**.

To delete all of the SMTP Routes, select the check box labeled "All" and click **Delete**.

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.

Procedure

Step 1 Click **Export SMTP Routes** on the SMTP Routes page.

Step 2 Enter a name for the file and click **Submit**.

Importing SMTP Routes

Similar to the Host Access Table (HAT) and the Recipient Access Table (RAT), you can also modify SMTP routes mappings by exporting and importing a file.

Procedure

Step 1 Click **Import SMTP Routes** on the SMTP Routes page.

Step 2 Select the file that contains the exported SMTP Routes.

Step 3 Click **Submit**. You are warned that importing will replace all existing SMTP Routes. All of the SMTP Routes in the text file are imported.

Step 4 Click **Import**.

You can place "comments" in the file. Lines that begin with a '#' character are considered comments and are ignored by AsyncOS. For example:

```
# this is a comment, but the next line is not
ALL:
```

At this point, our Email Gateway configuration looks like this:

Figure 11-1 SMTP Routes Defined for a Public Listener

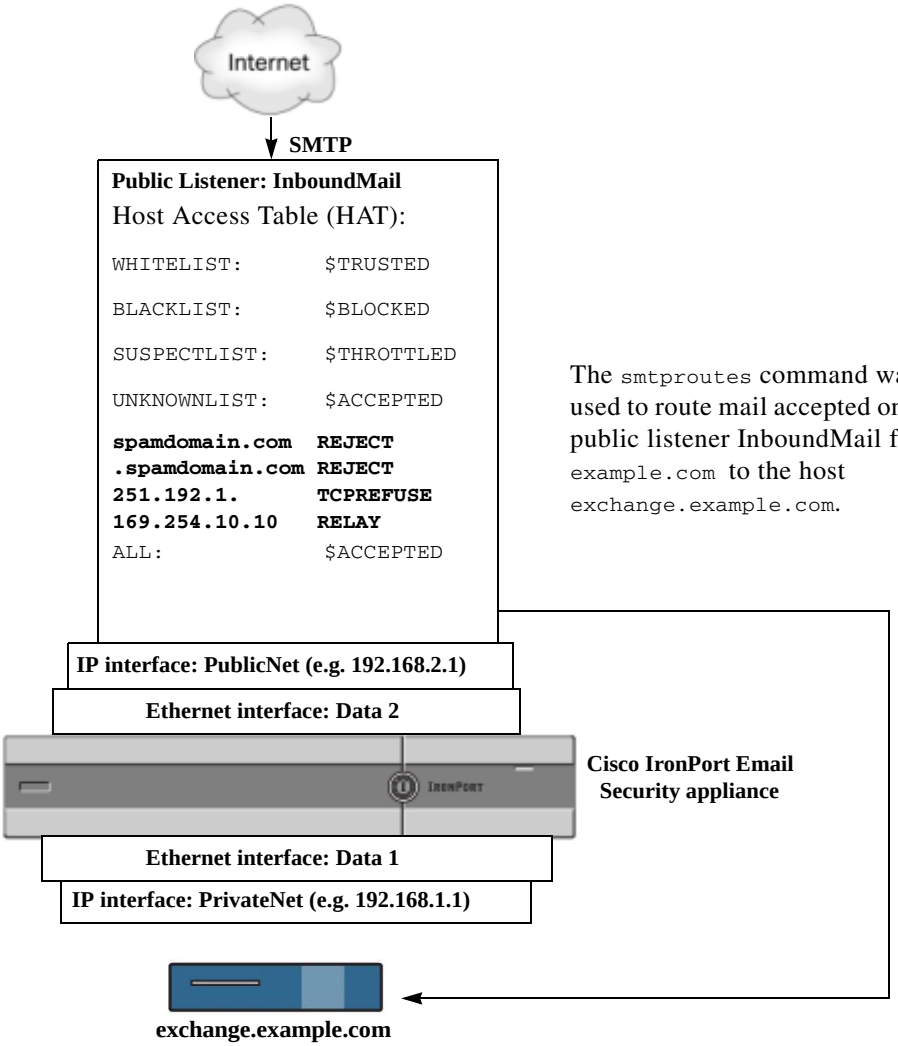

The `smtproutes` command was used to route mail accepted on the public listener InboundMail for `example.com` to the host `exchange.example.com`.

CHAPTER 12

Distributing Administrative Tasks

You can distribute administrative tasks for managing the Security Management appliance, and control access to the appliance in a variety of ways.

- About Distributing Administrative Tasks, page 12-1
- Assigning User Roles, page 12-1
- Managing Authentication of Administrative Users, page 12-9
- Additional Controls on Access to the Security Management Appliance, page 12-18
- Controlling Access to Sensitive DLP Information in Message Tracking, page 12-21
- Viewing Administrative User Activity, page 12-21

About Distributing Administrative Tasks

You can distribute administrative tasks on the 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

The Security Management appliance offers both predefined and custom user roles for monitoring and managing your email and web security appliances.

- Predefined User Roles, page 12-2
- Custom User Roles, page 12-4

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 12-1 Descriptions of User Roles

User Role Name	Description	Web Reporting/ 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/Yes
Administrator	User accounts with the Administrator role have full access to all configuration settings of the system.	Yes/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. Otherwise, they have the same privileges as the Administrator role.	Yes/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 the System Capacity report under the Email tab
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 the spam quarantine, if access is enabled. Users with this role cannot access the file system, FTP, or SCP.	Yes/No

Table 12-1 Descriptions of User Roles

User Role Name	Description	Web Reporting/ Scheduled Reports Capability
Guest	Users accounts with the Guest role can view status information and manage messages in the spam quarantine, if access is enabled. Users with the Guest role cannot access Message Tracking.	Yes/No
Web Administrator	User accounts with the Web Administrator role have access to all configuration settings under the **Web** tab.	Yes/Yes
Web Policy Administrator	User accounts with the Web Policy Administrator role can access the Web Appliance Status page and all pages in the Configuration Master. The web policy administrator can configure identities, access policies, decryption policies, routing policies, proxy bypass, custom URL categories, and time ranges. The web policy administrator cannot publish configurations.	No/No
URL Filtering Administrator	User accounts with the URL Filtering Administrator role can configure URL filtering only.	No/No
Email Administrator	User accounts with the Email Administrator role have access to all configuration settings within the Email menu only, including the spam quarantine.	No/No
Help Desk User	User accounts with the Help Desk User role are restricted to: • Message Tracking • Managing messages in the spam quarantine 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 the spam quarantine to allow access by this user.	No/No
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 or a new Custom Web 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, page 12-4.	No/No

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, Web Administrator, Web Policy Administrator, URL Filtering Administrator, and custom user.

If you use an LDAP directory to authenticate users, you assign directory groups to user roles instead of individual users. When you assign a directory group to a user role, each user in that group receives the permissions defined for the user role. For more information, see External User Authentication, page 12-16.

Before users can access the spam quarantine, you must enable that access. See Configuring Administrative User Access to the Cisco IronPort Spam Quarantine, page 7-7.

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, page 12-4
- About Custom Web User Roles, page 12-7

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
- Spam quarantine

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 *Cisco IronPort AsyncOS for Email Security Daily Management Guide*.

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.

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:

- Overview
- Incoming Mail
- Outgoing Destinations
- Outgoing Senders
- Internal Users
- DLP Incidents
- Content Filters
- Virus Types
- TLS Connections
- Outbreak Filters
- System Capacity
- Reporting Data Availability
- 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:

- Overview
- Incoming Mail
- Outgoing Destinations
- Outgoing Senders
- Internal Users
- Content Filters
- Virus Types
- Outbreak Filters
- Reporting Data Availability
- 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

- Reporting Data Availability
- Archived Reports

Message Tracking

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 DLP Information in Message Tracking, page 12-21.

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 Email Messages".

Quarantines

If you grant a custom role access to quarantines, users to whom you assign this role can search for, view, release, or delete messages the Spam Quarantine on this Security Management appliance.

Before users can access the spam quarantine, you must enable that access. See Configuring Administrative User Access to the Cisco IronPort Spam Quarantine, page 7-7,.

For complete information about the spam quarantine, including instructions for setting up your appliances to enable access to this quarantine from the Security Management appliance, see the "Managing the Cisco IronPort Spam Quarantine" chapter.

Creating Custom Email User Roles

You can create custom email user roles for access to Email Reporting, Message Tracking, and the spam quarantine.

For descriptions of the access that each of these options permits, see About Custom Email User Roles 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.

Procedure

Step 1 Choose **Management Appliance > System Administration > User Roles**.

Step 2 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 3 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.

Step 4 Choose the access privileges to enable for this role.

Step 5 Click **Submit** to return to the User Roles page, which lists the new user role.

Step 6 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 7 Commit your changes.

Step 8 If you granted this role access to the spam quarantine, enable access for this role:

See:

- Configuring Administrative User Access to the Cisco IronPort Spam Quarantine, page 7-7.

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 12-1 Account Privileges Page for a Delegated Administrator assigned Custom Email User Roles

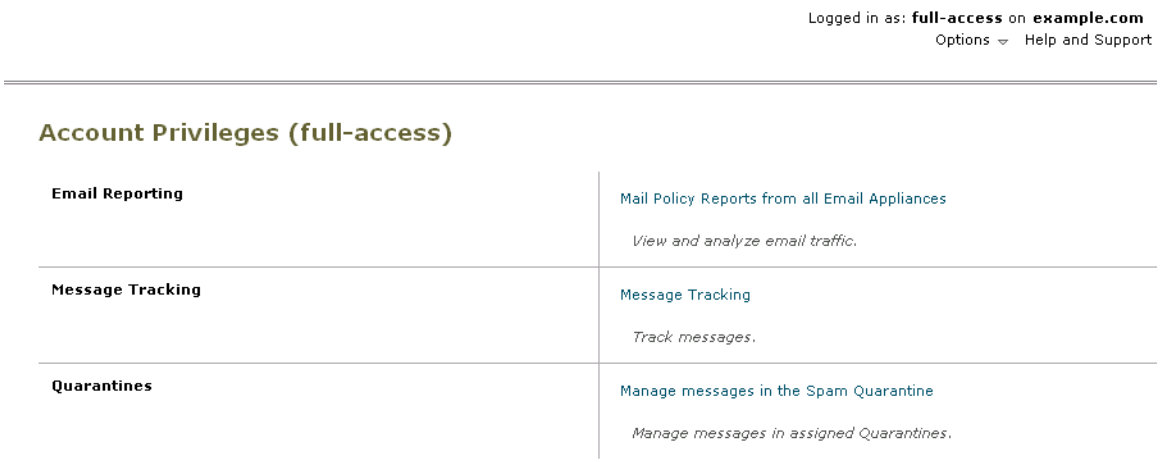

About Custom Web User Roles

Custom web user roles allow users to publish policies to different Web Security appliances, and gives them the permission to edit or publish the custom configuration to different appliances.

From the **Web > Configuration Master > Custom URL Categories** page on the Security Management appliance, you can view the URL categories and policies that you are allowed to administer and publish. Additionally, you can go to the **Web > Utilities > Publish Configuration Now** page and view the possible configurations.

Note Remember that when you create a custom role with Publish Privilege capabilities, when user logs in, they will not have any usable menus. They do not have the publish menu and they will land on an non-editable landing screen since the URL and policy tabs do not have any capabilities. In effect, you have a user that cannot publish or administer any categories or policies.

The workaround to this issue is that if you want a user to be able to publish, but not to be able to manage any categories or policies, you **must** create a custom category which is not used in any policy, and give that user the ability to manage that custom category along with publishing. In this way, if they add or delete URLs from that category, it does not affect anything.

You can delegate web administration by creating and editing custom user roles.

- Creating Custom Web User Roles
- Editing Custom Web User Roles

Creating Custom Web User Roles

Procedure

Step 1 Choose **Management Appliance > System Administration > User Roles**.

Step 2 Click **Add Web User Role**.

Tip Alternatively, you can create a new role by duplicating an existing Web User Role: Click the Duplicate icon in the applicable table row, then modify the resulting copy.

Step 3 Enter a unique name for the user role (for example, “canadian-admins”) and a description.

Note The name must contain only lowercase letters, numbers, and dashes. It cannot start with a dash.

Step 4 Choose whether you want the policies and custom URL categories to be visible or hidden by default.

Step 5 Choose whether you want Publish privileges turned on or off.

This privilege allows the user to publish any Configuration Master for which the user can edit Access Policies or URL Categories.

Step 6 Choose whether to start with new (empty) settings or to copy an existing custom user role. If you choose to copy an existing user role, choose from the list the role that you want to copy.

Step 7 Click **Submit** to return to the User Roles page, which lists the new user role.

Note If you have enabled the anonymized feature within web reporting, all user roles with access to web reporting will have unrecognizable user names and roles in the interactive reports page. See the Scheduling Web Reports section in Chapter 5, “Using Centralized Web Reporting and Tracking.” The exception is the Administrator role, which is able to see actual user names in the scheduled reports. If the anonymize feature is enabled, scheduled reports that are generated by the Operator and Web Administrator are anonymized.

Note If you use the **Web > Utilities > Security Services Display > Edit Security Services Display** page to hide one of the Configuration Masters, the User Roles page also hides the corresponding Configuration Master column; however, privilege settings for the hidden Configuration Master are retained.

Editing Custom Web User Roles

Procedure

Step 1 On the User Roles page, click the role name to display the Edit User Role page.

Step 2 Edit any of the settings: name, description, and visibility of policies and custom URL categories.

Step 3 Click **Submit.**

To edit privileges for a custom user role:

Navigate to the User Roles page.

- To edit access policy privileges, click "Access policies" to display a list of access policies configured in the Configuration Master. In the Include column, select the check boxes of the policies to which you want to give the user edit access. Click **Submit** to return to the User Roles page.

-or-

- To edit custom URL category privileges, click Custom URL Categories to display a list of the custom URL categories defined on the Configuration Master. In the Include column, select the check boxes of the custom URL categories to which you want to give the user edit access. Click **Submit** to return to the User Roles page.

Managing Authentication of Administrative Users

You can control access to the appliance by defining authorized users locally on the appliance, and/or by using external authentication.

- Managing Locally-Defined Administrative Users, page 12-9
- External User Authentication, page 12-16

Managing Locally-Defined Administrative Users

- Adding Locally-Defined Users, page 12-10
- Editing Locally-Defined Users, page 12-10
- Deleting Locally-Defined Users, page 12-11
- Viewing the List of Locally-Defined Users, page 12-11
- Setting and Changing Passwords, page 12-12

- Configuring Restrictive User Account and Password Settings, page 12-12
- Requiring Users to Change Passwords at Next Login, page 12-14
- Locking and Unlocking Local User Accounts, page 12-15

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.

Procedure

Step 1 If you will assign custom user roles, we recommend that you define those roles first. See Custom User Roles, page 12-4.

Step 2 On the Security Management appliance, choose **Management Appliance > System Administration > Users**.

Step 3 Click **Add User**.

Step 4 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 5 Enter a full name for the user.

Step 6 Select a predefined role or a custom role. See Table 12-1 for more information about user roles.

If you add a new Email role or Web role here, enter a name for the role. For naming restrictions, see Creating Custom Email User Roles, page 12-6 or Creating Custom Web User Roles, page 12-8.

Step 7 Enter a password and reenter it. Passwords must be at least six characters long.

Step 8 Submit and commit your changes.

Step 9 If you added a custom user role on this page, assign privileges to that role now. See Custom User Roles, page 12-4.

Editing Locally-Defined Users

Use this procedure to change a password, for example.

Procedure

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

Procedure

Step 1 Click the trash can icon corresponding to the user's name in the Users listing.

Step 2 Confirm the deletion by clicking **Delete** in the warning dialog that appears.

Step 3 Click **Commit** to commit your changes.

Viewing the List of Locally-Defined Users

To view a list of locally-defined users, choose **Management Appliance > System Administration > Users**.

Figure 12-2 Example Users Page

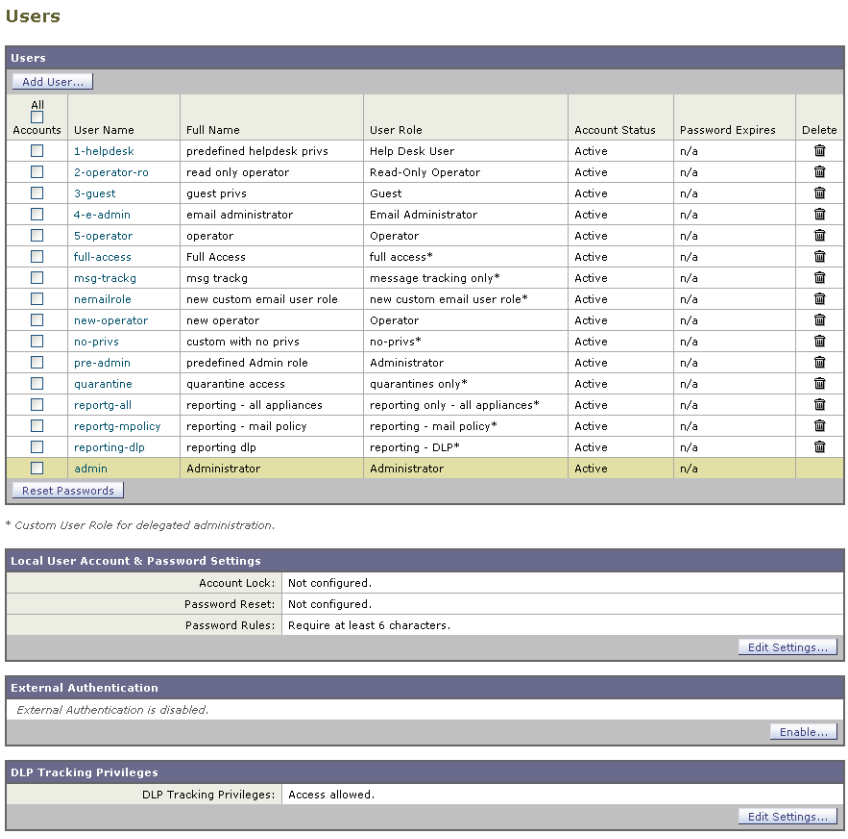

Note Asterisks indicate users assigned custom user roles for delegated administration. "Unassigned" appears in red if the user's custom role has been deleted. For more information on custom user roles, see Custom User Roles, page 12-4.

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, page 12-10 for more information).
- To change the password for the default admin user account for the system, use the Edit User page in the GUI (see Editing Locally-Defined Users, page 12-10 for more information) or use the `password` or `passwd` command in the CLI. If you forget the password for the "admin" user account, contact your customer support provider to reset the password.
- To force users to change their passwords, see Requiring Users to Change Passwords at Next Login, page 12-14.
- 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.

  On the Change Password page, enter the old password, and then enter the new password and reenter it for confirmation. Click **Submit** to log out. The login screen appears.

Configuring Restrictive User Account and Password Settings

You can define user account and password restrictions to enforce organizational password policies. The user account and password restrictions apply to local users defined on the Cisco IronPort 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.

Procedure

Step 1 Choose **Management Appliance > System Administration > Users**.

Step 2 Scroll down to the **Local User Account and Password Settings** section.

Step 3 Click **Edit Settings.**

Step 4 Configure the settings described in Table 12-2.

Table 12-2 Local User Account and Password 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, page 12-15.
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 at Next Login, page 12-14. 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 uses. After a password expires, the user is forced to change the account password at the next login. **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.
Password Rules: Require at least <number> characters.	Enter the minimum number of characters that passwords may contain. You can enter any number from 6 to 128. Default is six (6).
Password Rules: Require at least one number (0-9).	Choose whether or not the passwords must contain at least one number.

Table 12-2 Local User Account and Password Settings (continued)

Setting	Description

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).

Step 5 Submit and commit your changes.

Requiring Users to Change Passwords at Next Login

To require all or selected users to change their passwords the next time they access the Security Management appliance, 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 Configuring Restrictive User Account and Password Settings, page 12-12.

Procedure

Step 1 Choose **Management Appliance > System Administration > Users**.

Step 2 In the Users section, select the check boxes beside the users who will be required to change passwords on next login.

Step 3 Click **Reset Passwords**.

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 Configuring Restrictive User Account and Password Settings, page 12-12.
- Administrators can manually lock user accounts. See Locking User Accounts Manually, page 12-15.

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

Procedure

Step 1 First time only: Set up the appliance to enable user account locking:

a. Go to **Management Appliance > System Administration > Users**.

b. In the **Local User Account & Password Settings** section, click **Edit Settings**.

c. Select the checkbox to **Display Locked Account Message if Administrator has manually locked a user account** and enter your message.

d. Submit the change.

Step 2 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, page 12-15.

Step 3 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 *Cisco IronPort AsyncOS for Email Configuration Guide* 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 Cisco IronPort appliance to use the external directory to authenticate users who log in to the Security Management appliance.

Note

- Some features described in Customizing Your View, page 13-52 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.

Configuring LDAP Authentication

To configure LDAP authentication, see Configuring External Authentication of Administrative Users Using LDAP, page 10-14.

Enabling RADIUS Authentication

You can use a RADIUS directory to authenticate users and assign groups of users to Cisco IronPort roles. The RADIUS server should support the CLASS attribute, which AsyncOS uses to assign users in the RADIUS directory to Cisco IronPort 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.

Procedure

Step 1 On the Management Appliance > System Administration > Users page, click **Enable**.

Step 2 Select the **Enable External Authentication** check box.

Step 3 Select RADIUS for the authentication type.

Step 4 Enter the host name for the RADIUS server.

Step 5 Enter the port number for the RADIUS server. The default port number is 1812.

Step 6 Enter the Shared Secret password for the RADIUS server.

Note When enabling external authentication for a cluster of Cisco IronPort appliances, enter the same Shared Secret password on all appliances in the cluster.

Step 7 Enter the number of seconds that the appliance waits for a response from the server before timing out.

Step 8 Select whether to use Password Authentication Protocol (PAP) or Challenge Handshake Authentication Protocol (CHAP) for the authentication protocol.

Step 9 (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 10 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 11 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 • Web Administrator • Web Policy Administrator • URL Filtering Administrator • Technician • Operator • Cloud Administrator • 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 12 (Optional) Click **Add Row** to add another group. Repeat step 11 for each group of users that the appliance authenticates.

Step 13 Submit and commit your changes.

Additional Controls on Access to the Security Management Appliance

- Configuring IP-Based Network Access, page 12-18
- Configuring the Web UI Session Timeout, page 12-20

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. Figure 12-3 shows the Network Access page with a list of user IP addresses that are allowed to connect directly to the Security Management appliance.

Figure 12-3 Example Network Access Settings

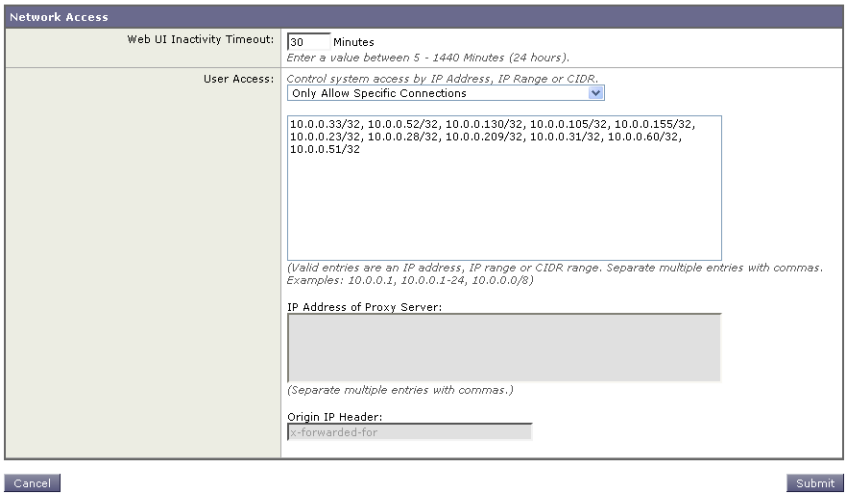

AsyncOS offers four different modes of control for the access list:

- **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.

Procedure

Step 1 Choose **System Administration > Network Access**.

Step 2 Click **Edit Settings**.

Step 3 Select the mode of control for the access list.

Step 4 Enter the IP addresses from which users will be allowed to connect to the appliance.

You can enter an IP address, IP address range or CIDR range. Use commas to separate multiple entries.

Step 5 If connecting through a proxy is allowed, enter the following information:

- The IP addresses of the proxies allowed to connect to the appliance. Use commas to separate multiple entries.
- The name of the origin IP header that the proxy sends to the appliance, which contains the IP addresses of the remote user's machine and the proxy servers that forwarded the request. By default, the name of the header is `x-forwarded-for`.

Step 6 Submit and commit your changes.

Configuring 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 IronPort Spam Quarantine sessions, which have a 30 minute timeout that cannot be configured.

Procedure

Step 1 Use the **System Administration > Network Access** page.

Step 2 Click **Edit Settings**.

Step 3 Enter the number of minutes users can be inactive before being logged out. You can define a timeout period between 5 and 1440 minutes.

Step 4 Submit and commit your changes.

Controlling Access to Sensitive DLP Information in Message Tracking

Messages that violate Data Loss Prevention (DLP) policies typically include sensitive information, such as corporate confidential information or personal information including credit card numbers or health records. By default, this content appears in the DLP Matched Content tab on the Message Details page for messages listed in Message Tracking results.

You can choose to hide this tab and its content from Security Management appliance users based on their assigned predefined or custom role:

Procedure

Step 1 Go to the **Management Appliance > System Administration > Users** page.

Step 2 In the **DLP Tracking Privileges** section, click **Edit Settings**.

Step 3 Select the roles for which you want to grant access to DLP data in Message Tracking.

Only custom roles with access to Message Tracking are listed.

Step 4 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.

Viewing Administrative User Activity

- Viewing Active Sessions Using the Web, page 12-21
- Viewing Administrative User Activity via the Command Line Interface, page 12-22

Viewing Active Sessions Using the Web

From the Security Management appliance, you can view all active sessions and users logged in to the appliance.

Procedure

Step 1 From the upper right corner of the window, choose **Options > Active Sessions**.

Figure 12-4 Active Sessions Menu

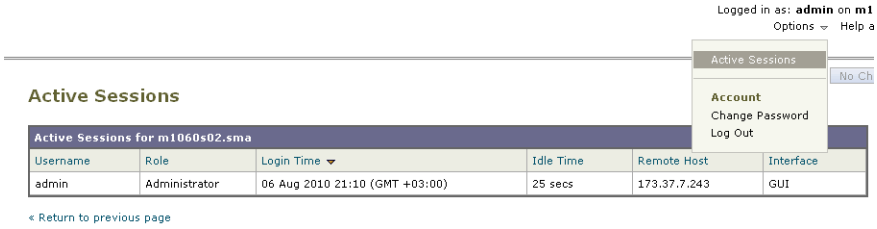

From the Active Sessions page you can view the User name, what role the user has, the time the user logged in, idle time, and whether the user is logging in from the command line or the GUI.

Viewing 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, the time of login, the idle time, and the remote host from which the user is logged in:

```
mail3.example.com> who

Username  Login Time  Idle Time  Remote Host  What
========  ==========  =========  ===========  ====
admin     03:27PM     0s         10.1.3.201   cli
```

- The `whoami` command displays the 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
```

CHAPTER 13

Common 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 9, "Monitoring System Status."

Note Several of the features or commands described in this chapter can affect routing precedence. For more information, see IP Addresses, Interfaces, and Routing, page B-3.

- Working with Feature Keys, page 13-2
- Performing Maintenance Tasks Using CLI Commands, page 13-3
- Backing Up Your Security Management Appliance, page 13-6
- Disaster Recovery on the Security Management Appliance, page 13-12
- Upgrading Appliance Hardware, page 13-14
- Upgrading AsyncOS, page 13-16
- About Reverting to an Earlier Version of AsyncOS, page 13-25
- About Updates, page 13-27
- Configuring the Return Address for Generated Messages, page 13-28
- Managing Alerts, page 13-28
- Changing Network Settings, page 13-36
- Configuring the System Time, page 13-41
- Saving and Importing Configuration Settings, page 13-43
- Managing Disk Usage, page 13-50
- Customizing Your View, page 13-52

Working with Feature Keys

On the main Security Management appliance, choose **Management Appliance > System Administration > Feature Keys** on the GUI (or the **featurekey** command from the command line prompt) to enter the key and enable the associated functionality.

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. If you incorrectly enter a key, an error message is generated.

Occasionally, Cisco IronPort Customer Support may provide a key to enable specific functionality on your system.

Two pages provide feature keys functionality: the Feature Keys page and the Feature Key Settings page.

Feature Keys Page

Log in to the Security Management appliance, and select **Management Appliance > System Administration > Feature Keys**. Use the Feature Keys page to perform the following tasks:

- 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.

The Feature Keys for Serial Number: <*Serial Number*> section lists the enabled features for the appliance. The Pending Activation section lists feature keys that have been issued for the appliance but have not yet been activated. By default, the appliance periodically checks for new keys. You can change the appliance configuration to modify this behavior. In addition, you can click the **Check for New Keys** button to refresh the list of pending keys.

Figure 13-1 Feature Keys Page

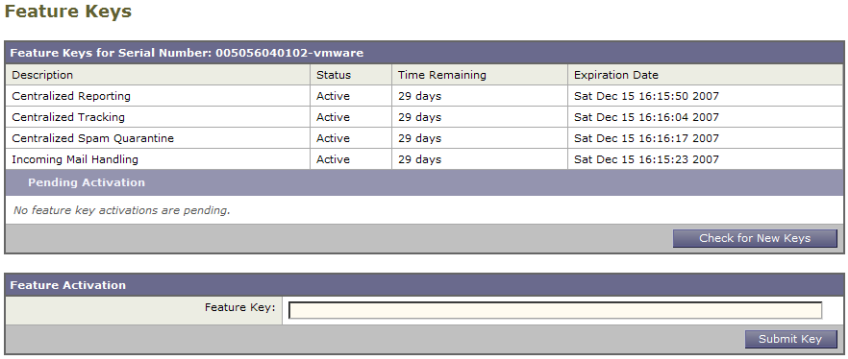

To add a new feature key manually, paste or enter the key into the Feature Key field and click **Submit Key.** An error message appears if the feature is not added (for example, if the key is incorrect); otherwise, the feature key is added to the list.

To activate a new feature key from the Pending Activation list, select the key (select the Select check box) and click **Activate Selected Keys**.

You can configure your appliance to automatically download and install new keys as they are issued. In this case, the Pending Activation list is always empty.

Feature Key Settings Page

Use the **Management Appliance > System Administration > Feature Key** Settings page to control whether the appliance checks for and downloads new feature keys, and whether or not the keys are automatically activated.

Expired Feature Keys

If the feature key for the feature you are trying to access has expired, contact your Cisco representative or other customer support organization.

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**
- **offline**
- **resume**
- **resetconfig**
- **version**

Shutting Down the Security Management Appliance

To shut down your Security Management appliance, 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 `reboot` 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.

Placing the Security Management Appliance into a Maintenance State

If you want to perform system maintenance, place the Security Management appliance into the offline state. The `suspend` and `offline` commands put AsyncOS into offline state. The offline state is characterized by the following:

- Inbound email connections are not accepted.
- Outbound email delivery is halted.
- Log transfers are halted.
- The CLI remains accessible.

You must enter a delay for the appliance to enter the offline state. 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, the offline state commences immediately.

Note The difference between the **suspend** command and the **offline** command is that the **suspend** command retains its state even after the machine is rebooted. If you issue the **suspend** command and reboot the appliance, you must use the **resume** command to return the system to an online state.

See also:

- "Suspending Email Delivery," "Resuming Email Delivery," "Suspending Receiving," and "Resuming Receiving" in the *Cisco IronPort AsyncOS for Email Security Advanced User Guide*

The suspend and offline Commands

```
mail3.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.

mail3.example.com> offline

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.
```

Resuming from an Offline State

The `resume` command returns AsyncOS to normal operating state after using the **suspenddel** or **suspend** command.

The resume Command

```
mail3.example.com> resume

Receiving resumed.
Mail delivery resumed.
mail3.example.com>
```

Resetting to Factory Defaults

When physically transferring the appliance, you may want to start with factory defaults. The Reset button in the Reset Configuration section of the **Management Appliance > System Administration > Configuration File** page, or the **resetconfig** command, resets *all* AsyncOS configuration values to factory defaults. This command is extremely destructive, and it should only be used when you are transferring the unit or as a last resort to solving configuration issues. It is recommended that you run the System Setup Wizard after resetting the configuration.

Note The `resetconfig` command works only when the appliance is in the offline state. When the **resetconfig** command completes, the appliance is automatically returned to the online state. If mail delivery was suspended before you issued the **resetconfig** command, the mail will attempt to be delivered again when the **resetconfig** command completes.

Warning **The `resetconfig` command will return all network settings to factory defaults, potentially disconnecting you from the CLI, disabling services that you used to connect to the appliance (FTP, Telnet, SSH, HTTP, HTTPS), and even removing additional user accounts you created with the `userconfig` command. Do not use this command if you are not able to reconnect to the CLI using the Serial interface or the default settings on the Management port through the default admin user account.**

The resetconfig Command

```
mail3.example.com> offline

Delay (seconds, minimum 30):
[30]> 45

Waiting for listeners to exit...
Receiving suspended.
Waiting for outgoing deliveries to finish...
Mail delivery suspended.

mail3.example.com> resetconfig

Are you sure you want to reset all configuration values? [N]> Y

All settings have been restored to the factory default.
```

Displaying the Version Information for AsyncOS

Procedure

Step 1 On the Security Management appliance, choose **Management Appliances > Centralized Services > System Status**.

Step 2 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.

Backing Up Your Security Management Appliance

- What Data Is Backed Up, page 13-6
- Restrictions and Requirements for Backups, page 13-6
- Backup Duration, page 13-7
- Availability of Services During Backups, page 13-8
- Interruption of a Backup Process, page 13-8
- Scheduling Single or Recurring Backups, page 13-9
- Starting an Immediate Backup, page 13-10
- Checking Backup Status, page 13-11
- Other Important Backup Tasks, page 13-12

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
- Email tracking (message tracking), including messages and meta data
- Web tracking
- Reporting (Email and Web)
- 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, page 13-12.

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	This release of AsyncOS for Security Management data can be backed up only to another Security Management appliance running the same release of AsyncOS for Security Management. If there is a version incompatibility, upgrade the target Security Management appliance 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 Chapter 2, "Setup, Installation, and Basic Configuration."
Communication between 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.
Appliance capacity	The 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 on the target appliance cannot be less than the corresponding allocation on the source appliance. **Note** 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 for all of the data being backed up. For example, if the source appliance is an M1060 and a smaller M650 is the target, you must reduce the space allocated on the larger M1060 to match the space available on the smaller M650 appliance. See Managing Disk Usage, page 13-50 for information on allocating disk space.
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 the service is brought back up for both the source and the target. 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 that a system backup is in progress:

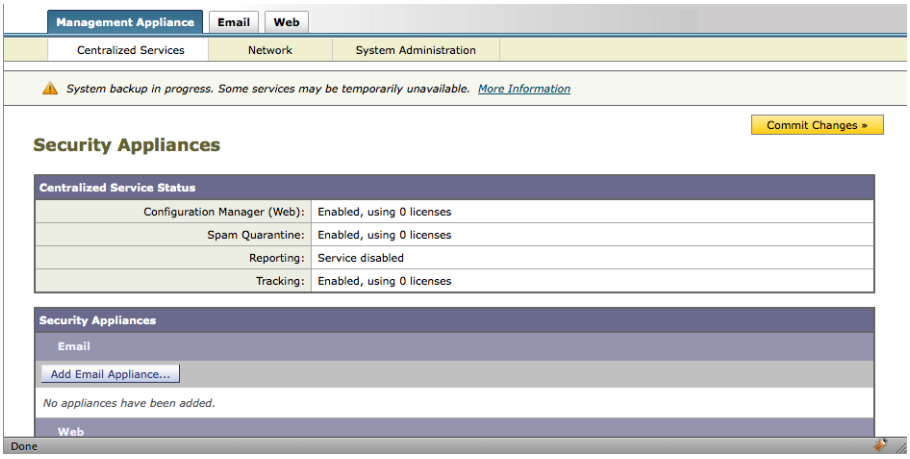

Interruption of a Backup Process

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.

Scheduling Single or Recurring Backups

You can schedule a single or recurring backup to occur at a predetermined time.

Before You Begin

Load a configuration file that matches the configuration of the source appliance onto the target appliance.

Note A backup process will not start if there are any ongoing backups on the remote machine.

Procedure

Step 1 Login, as administrator, to any SSH session.

Step 2 At the command prompt, type **backupconfig** and press **Enter**.

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

- `View – Allows you to view the scheduled backups`
- `Verify – Verifies whether the backup can be scheduled to a remote machine.`
- `Schedule - Schedule a backup to an appliance.`
- `Cancel – Cancels a scheduled backup.`
- `Status – Show the status of the ongoing back up in progress.`
- `Setup - Configure backup parameters.`

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.

The Security Management appliance now verifies the existence of the target machine, and determines whether the target machine has enough space to accept the data.

If the space on the target machine is not sufficient, the following error message is generated `"Backup cannot be scheduled. Reason: There is not enough space for Spam Quarantine, Email Tracking, Web Tracking, Reporting. Please increase disk allocation for these services on the target machine."` and data is not transferred.

Once the target machine is verified, the following choices appear:

```
1. Setup Repeating Backup Schedule—allows you schedule a periodic backup.
```

```
2. Schedule a single backup—allows you to schedule a single backup.

3. Start a Single Backup Now—allows you to initiate an immediate backup.
```

Step 9 To schedule a single backup, type 2 and press **Enter**.

Step 10 To schedule a recurring backup:

a. Type **1** and press **Enter**.

b. The following choices appear: `1. Daily, 2. Weekly, 3. Monthly.`

c. Choose the time frame 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, page 13-12.

Starting an Immediate Backup

Instant backups occur immediately when you initiate the **backupconfig** command on the CLI.

Before You Begin

Load a configuration file that matches the configuration of the source appliance onto the target appliance.

Note A backup process will not start if there are any ongoing backups on the remote machine.

Procedure

Step 1 Login, as administrator, to any SSH session.

Step 2 At the command prompt, type **backupconfig** and press **Enter**.

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

- `View – Allows you to view the scheduled backups.`
- `Verify – Verifies if the backup can be scheduled to a remote machine.`
- `Schedule – Allows you to schedule a backup to an appliance.`
- `Cancel – Cancels a scheduled backup.`
- `Status – Allows you to view the status of the ongoing back up in progress.`
- `Setup - Configure backup parameters.`

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.

The Security Management appliance now verifies the existence of the target machine, and determines whether the target machine has enough space to accept the data.

If the space on the target machine is not sufficient, the following error message is generated `"Backup cannot be scheduled. Reason: There is not enough space for Spam Quarantine, Email Tracking, Web Tracking, Reporting. Please increase disk allocation for these services on the target machine."` and data is not transferred.

Once the target machine is verified, the following choices appear on the console:

```
1. Setup Reoccurring Backup—allows you schedule a periodic backup.

2. Schedule a Single backup—allows you to schedule a single backups.

3. Start a Single Backup Now—allows you to initiate an immediate backup.
```

Step 9 Type **3** and press **Enter**.

Step 10 Enter a meaningful name for the backup job.

The backup process begins in a few minutes, and data starts transferring from the source machine to the target machine. After an immediate backup has started, the following message can be seen: `"Backup has been initiated and will begin in a few seconds."`

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, page 13-12.

Checking Backup Status

Checking Log Files

Backup logs record the backup process from start to finish.

Information about backup scheduling is in the SMA logs.

Checking Scheduled Backups

Procedure

Step 1 Login, as administrator, to any SSH session.

Step 2 At the command prompt, type **backupconfig** and press **Enter**.

Step 3 Choose the `View` operation.

Checking the Status of a Backup in Progress

Procedure

Step 1 Login, as administrator, to any SSH session.

Step 2 At the command prompt, type **backupconfig** and press **Enter**.

Step 3 Choose the `Status` operation.

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, page 13-43. 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, page 14-22.

  Additionally, you can set up a log subscription for Backup Logs. See Creating a Log Subscription in the GUI, page 14-23.

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 Your Security Management Appliance, page 13-6.

A typical appliance configuration might look like Figure 13-2:

Figure 13-2 Disaster Recovery: A Typical Environment

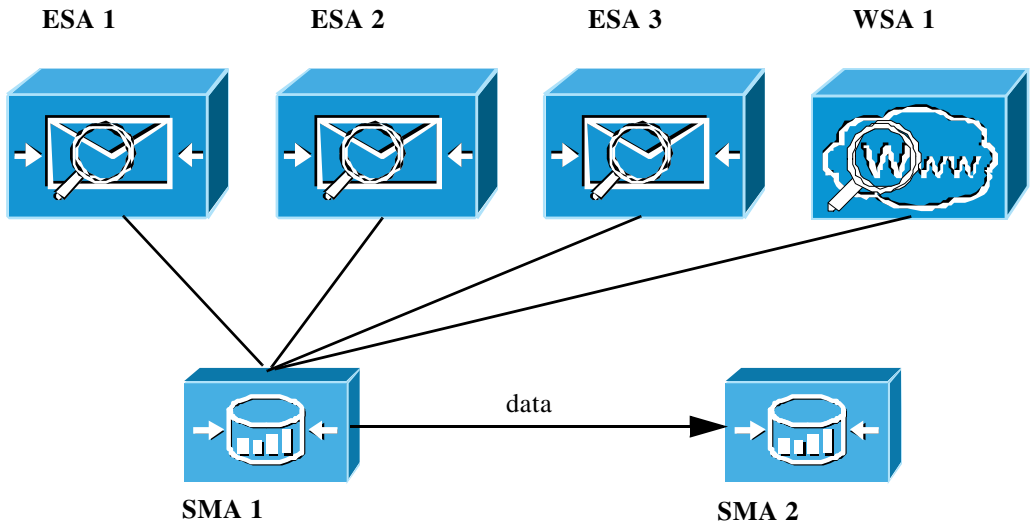

In this environment, SMA 1 is the primary Security Management appliance that is receiving data from ESAs 1-3 and WSA 1. SMA 2 is the backup Security Management appliance receiving backup data from SMA1.

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:

Procedure

Step 1 Load onto your backup Security Management appliance (SMA2) the configuration file that you saved from your primary Security Management appliance (SMA1).

For information, see Loading a Configuration File, page 13-45 or use the `loadconfig` command.

Step 2 Recreate the IP address from the failed SMA 1 to be the IP address on SMA 2:

a. On SMA 2 choose **Network > IP Interfaces > Add IP Interfaces**.

b. 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 on Adding IP Interfaces, see Configuring IP Interfaces, page A-2.

Step 3 Click **Submit** and **Commit**.

Step 4 Enable all of the services on the new Security Management appliance (SMA 2).

In this case, you must re-enable services for ESAs 1-3 and WSA 1. For information, see Configuring Services on the Security Management Appliance, page 2-14.

Step 5 Add all appliances on to the new Security Management appliance (SMA 2).

For information, see About Adding Managed Appliances, page 2-12.

Step 6 Test to see that each appliance is enabled and working by establishing a connection to the appliances and testing the connections.

SMA 2 is now the primary Security Management appliance. All data from ESAs 1-3 and WSA 1 now goes to SMA 2, as shown in Figure 13-3:

Figure 13-3 Disaster Recovery: Final Result

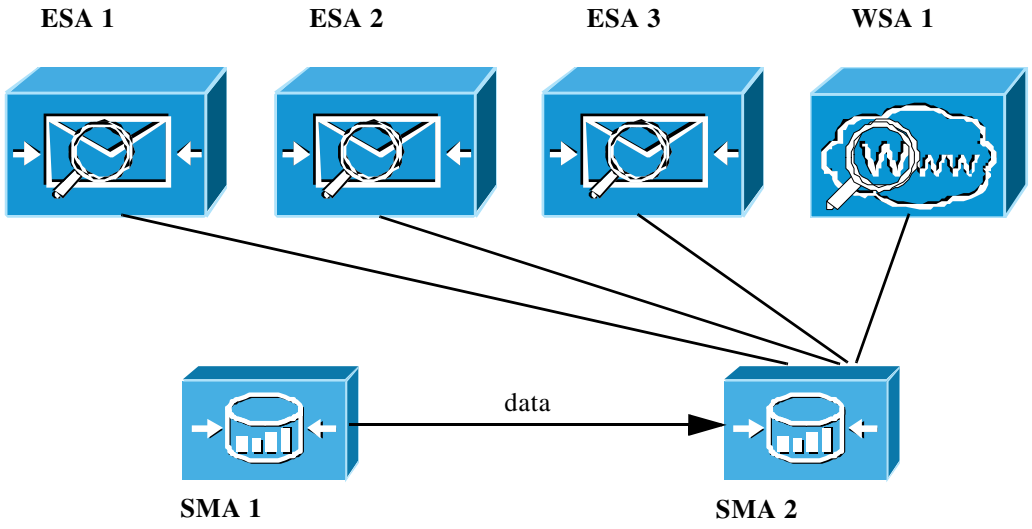

Restoring Other Data

See Other Important Backup Tasks, page 13-12 for additional data that you may need to restore.

Upgrading Appliance Hardware

If you are upgrading from an older Security Management appliance to a newer model, perform the following steps to transfer the data from the older appliance to the new appliance successfully.

Note It is important to remember that while it is possible to transfer data between Security Management appliances that are different sizes, the new appliance must have the same size allocations or greater.

Figure 13-4 Upgrading to a New Security Management Appliance Hardware

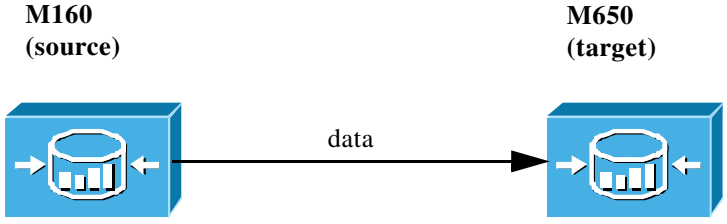

Before You Begin

- Understand the information in Backing Up Your Security Management Appliance, page 13-6.
- Meet the prerequisites described in Restrictions and Requirements for Backups, page 13-6.
- Save a copy of the configuration file from your source appliance to a location that you can reach from the target appliance. See Saving and Importing Configuration Settings, page 13-43.
- You may need to edit the configuration file before importing it into the new appliance. For example, change interface IP addresses to ensure that they are unique on the network, or make the interface names on the new appliance match the corresponding interface names on the old appliance.

Procedure

Step 1 Run the System Setup Wizard on the new appliance.

Step 2 Configure the appliance or import a configuration file.

Step 3 Login, as administrator, to an SSH session.

Step 4 At the command prompt, type **backupconfig** and press **Enter**.

Choose the operation you want to perform:

- View—Allows you to view the scheduled backups
- Verify—Verifies whether the backup can be scheduled to a remote machine.
- Schedule—Allows you to schedule a backup to an appliance.
- Cancel—Cancels a scheduled backup.
- Status—Allows you to view the status of the ongoing back up in progress.

Step 5 Type **Schedule** and press **Enter**.

Step 6 Type the IP address and name of the target Security Management appliance.

The Security Management appliance now verifies the existence of the target machine, and if the target machine has enough space to accept the data.

It is possible to transfer data between Security Management appliances that are different sizes, but, the new appliance must have the same size allocations or greater. If the space on the target machine is not sufficient, the following error message is generated **"Backup cannot be scheduled. Reason: There is not enough space for isq, tracking, reporting, slbl. Please increase disk allocation for these services on the target machine"** and data is not transferred.

Once the target machine is verified, the following choices appear on the console:

- 1. Setup Reoccurring Backup—allows you schedule a periodic backup.
- 2. Schedule a Single backup—allows you to schedule a single backups.
- 3. Start a Single Backup Now—allows you to initiate an immediate backup.

Step 7 Type **3** and press **Enter**.

The backup process begins, and data instantly starts transferring from the source machine to the target machine. After an immediate backup has started, the following message can be seen: `"Backup has been initiated and will begin in a few seconds."`

Step 8 Suspend all data transfer between the source appliance and the new target appliance by typing the **suspendtransfers** command at the command line prompt.

The **suspendtransfers** command stops the older source Security Management appliance from receiving any data.

Step 9 Run a new instant backup on the source machine by repeating steps 2 through 5 above.

What To Do Next

Use the `resumetransfers` command to restart data transfer.

To check the status of data transfers, go to **Centralized Services > System Status**.

Upgrading AsyncOS

This section discusses the following topics related to software upgrades and updates on the Security Management appliance:

- Batch Commands for Upgrades, page 13-16
- Determining Network Requirements, page 13-16
- Choosing an Upgrade Method: Remote vs. Streaming, page 13-16
- Configuring Upgrade and Service Update Settings, page 13-19
- Before You Upgrade: Important Steps, page 13-24
- Upgrading AsyncOS, page 13-24
- After Upgrading, page 13-25

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/en/US/products/ps10154/prod_command_reference_list.html.

Determining Network Requirements

Cisco IronPort uses a distributed upgrade server architecture to make sure customers can quickly download AsyncOS upgrades wherever in the world they are located. Because of this distributed server architecture, the Cisco IronPort update servers 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 IronPort 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 IronPort provides two methods (or 'sources') for upgrading AsyncOS on your Cisco IronPort appliances:

- Streaming upgrades — Each Cisco IronPort appliance downloads the AsyncOS upgrades via HTTP directly from the Cisco IronPort update servers.

- Remote upgrades — You only download the upgrade image from Cisco IronPort one time, and then serve it to your Cisco IronPort appliances. Your Cisco IronPort appliances download the AsyncOS upgrades from a server within your network.

You will configure the upgrade method in Configuring Upgrade and Service Update Settings, page 13-19. Optionally, use the **updateconfig** command in the CLI.

Streaming Upgrade Overview

In Streaming upgrades, each Cisco IronPort appliance connects directly to the Cisco IronPort update servers to find and download upgrades:

Figure 13-5 Streaming Update Method

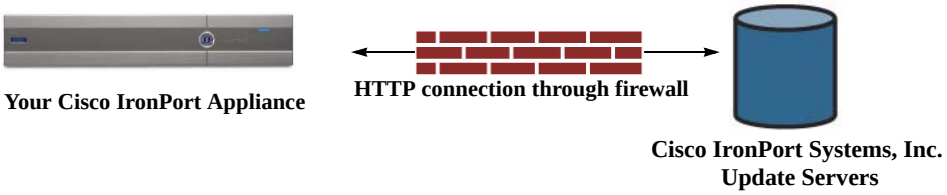

This method requires that your Cisco IronPort appliance contacts the Cisco IronPort Systems update servers directly from the network.

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 Cisco IronPort's 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 Cisco IronPort appliances.

Figure 13-6 Remote Update Method

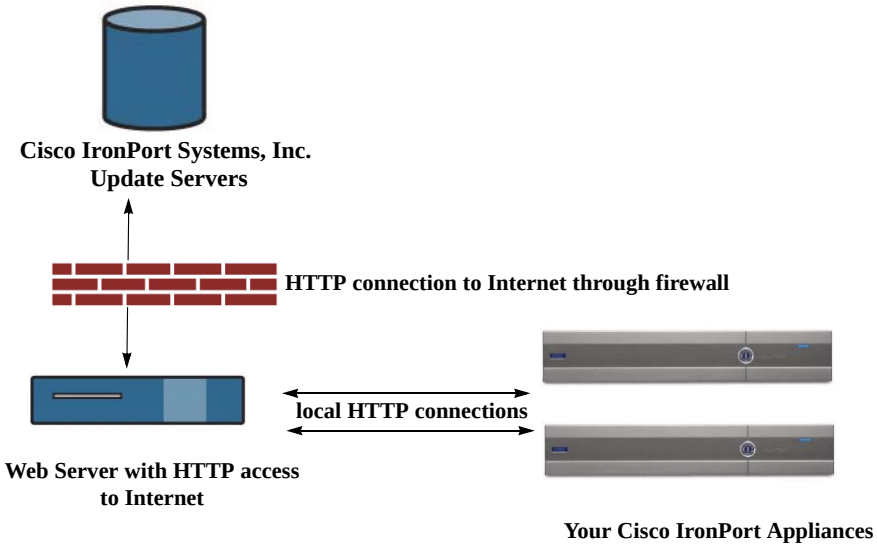

The basic process is as follows:

Procedure

Step 1 Read the information in Hardware and Software Requirements for Remote Upgrades, page 13-18 and Hosting a Remote Upgrade Image, page 13-18.

Step 2 Configure a local server to retrieve and serve the upgrade files.

Step 3 Download the upgrade files.

Step 4 Choose **Management Appliance > System Administration > Update Settings**

From this page, choose to configure the appliance to use the local server.

Step 5 Choose **Management Appliance > System Administration > System Upgrade**

Step 6 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, page 13-16.

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 Cisco IronPort Systems update servers.
- 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 IronPort 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 IronPort appliances on your network to the version selected at `http://updates.ironport.com/fetch_manifest.html`. Your Cisco IronPort appliances still download the upgrade from the Cisco IronPort Systems update 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, see the Cisco IronPort Knowledge Base or contact your Cisco IronPort 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 upgrading 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 Security Management 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 IronPort 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.

Upgrade and Update Settings

Table 13-1 describes the update and upgrade settings you can configure.

Table 13-1 Update Settings for Security Services

Setting	Description
Update Servers (images)	Choose whether to download Cisco IronPort AsyncOS upgrade and service update software images, such as time zone rules and Feature Key updates, from the Cisco IronPort update servers or a from a local web server. The default is the Cisco IronPort update 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, page 13-21. • You want to download Cisco IronPort AsyncOS upgrade images to your appliance at your convenience. (You can still download service update images dynamically from the Cisco IronPort 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, page 13-16 and Remote Upgrade Overview, page 13-17.
Update Servers (lists)	Choose whether to download the lists of available upgrades and service updates (the manifest XML files) from the Cisco IronPort update servers or from a local web server. The default for both upgrades and updates is the Cisco IronPort update 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, page 13-21. 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, page 13-16 and Remote Upgrade Overview, page 13-17.
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.

Table 13-1 Update Settings for Security Services (continued)

Setting	Description
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.

Static Upgrade and Update Server Settings for Environments with Strict Firewall Policies

The Cisco IronPort 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 13-7 Static URLs for Update Servers (images) Settings

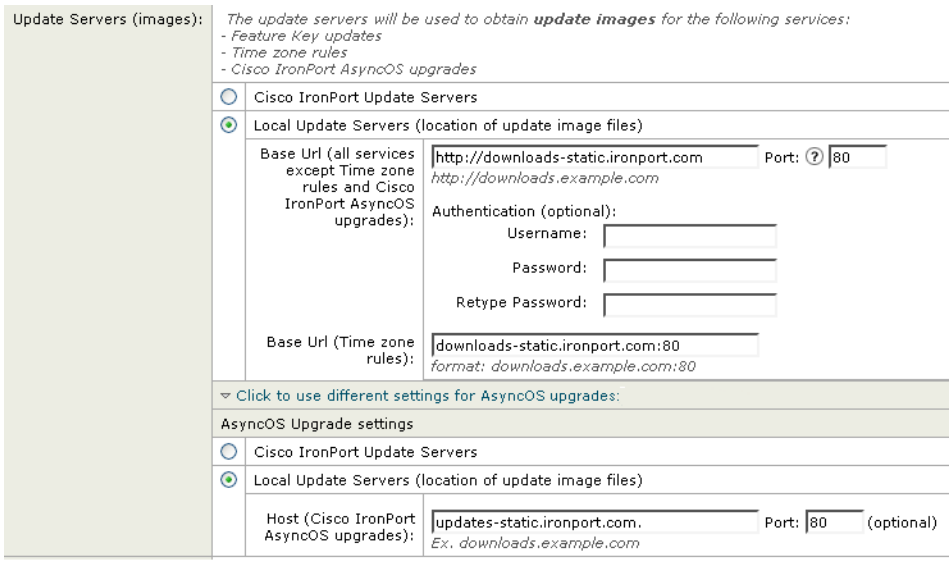

Figure 13-8 Static URLs for Update Servers (list) Settings

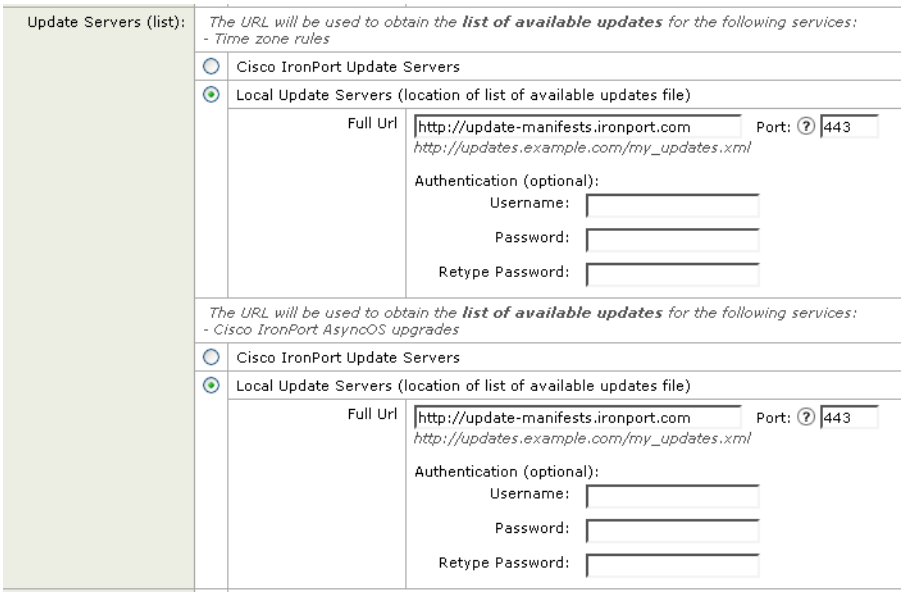

Table 13-2 Static Addresses for Environments with Strict Firewall Policies

Section	Setting	Static URL/IP Address and Port
Update Servers (images):	Base Url (all services except Time zone rules and Cisco IronPort 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 (Cisco IronPort AsyncOS upgrades)	updates-static.ironport.com 208.90.58.25 Port 80
Update Servers (list):	For updates: Full Url	update-manifests.ironport.com 208.90.58.5 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

Procedure

Step 1 Choose **Management Appliance > System Administration > Update Settings.**

Step 2 Click **Edit Update Settings**.

Use the descriptions in Upgrade and Update Settings, page 13-20 to configure the settings in this procedure.

Step 3 In the **Update Servers (images)** section, specify the servers from which to download images for updates.

Figure 13-9 Server Settings for Update Images

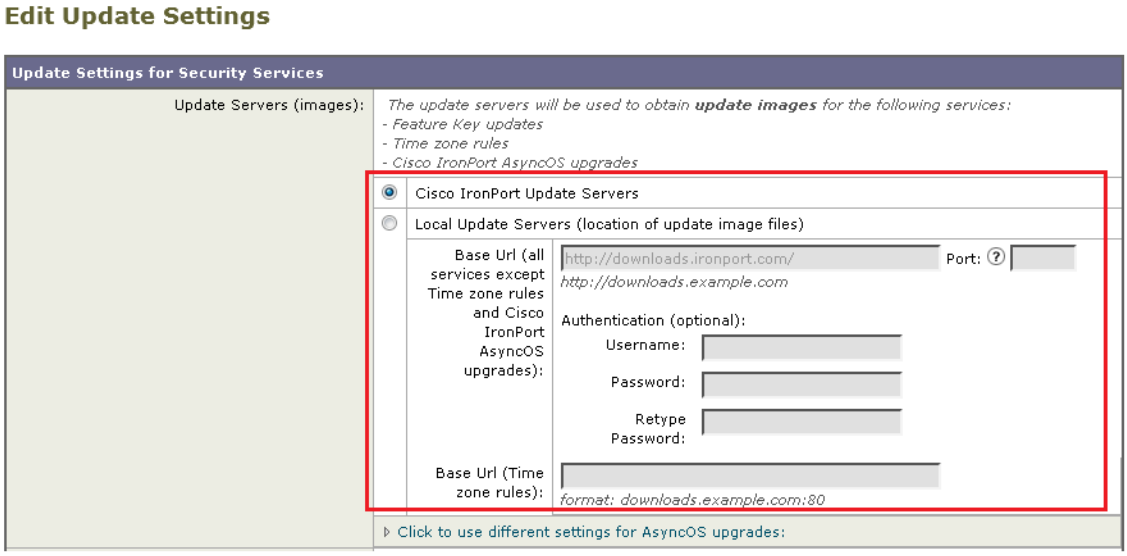

Step 4 Specify the server from which to download images for AsyncOS upgrades:

a. At the bottom of the same section, click the **Click to use different settings for AsyncOS upgrades** link:

Figure 13-10 Link to Specify Server Settings for Upgrade Images

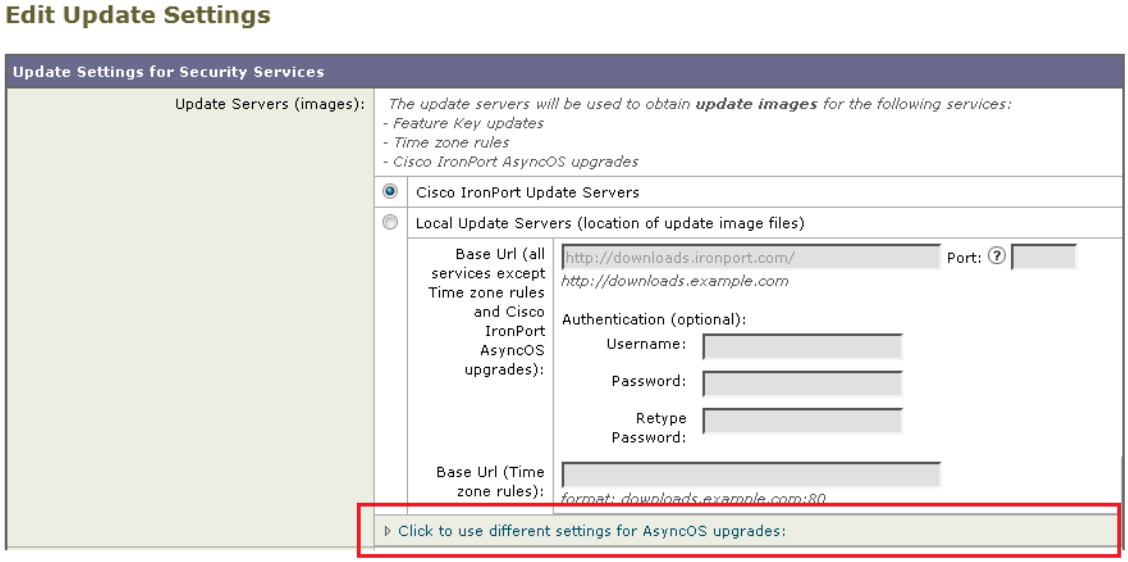

b. Specify server settings for downloading images for AsyncOS upgrades.

Step 5 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 6 Specify settings for Time Zone rules and interface.

Step 7 (Optional) Specify settings for Proxy Servers.

Step 8 Submit and commit your changes.

Step 9 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.

Note If you chose to download update images from the Cisco IronPort update servers, the "Feature Key updates" line on this page shows the URL of the Cisco IronPort update server, instead of "Dynamic (Cisco IronPort Update Server)."

Before You Upgrade: Important Steps

Procedure

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 Disk Space Maximums and Allocations, page 13-50.
- 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.

Step 3 If you are using the Safelist/Blocklist feature, export the list off the appliance.

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, page 13-19.

Upgrading AsyncOS

Before You Begin

Make sure you have completed the tasks in Before You Upgrade: Important Steps, page 13-24.

Procedure

Step 1 Choose **Management Appliance > System Administration > System Upgrade**.

Step 2 Click **Available Upgrades**.

Step 3 Select an upgrade from the list of available upgrades.

Step 4 In the Upgrade Preparation section, check the checkbox if you want to save the current configuration to the configuration directory before upgrading. This is recommended. Additionally, you can email this password file to an email of your choice by entering the email address in the text field.

In this section you can also hide the passwords in the configuration file by checking Mask Passwords in the Configuration File checkbox.

Step 5 Click **Begin Upgrade**. A progress bar appears near the top of the page. You may be asked one or more times to confirm changes or read and agree to new license agreements.

Step 6 Click **Continue** to complete the upgrade.

Step 7 After the upgrade is finished, you are prompted to reboot the appliance.

Step 8 Click **Reboot Now.**

What To Do Next

See additional recommended tasks in After Upgrading, page 13-25.

After Upgrading

After the upgrade is complete, complete the following:

- (For deployments with associated Email Security appliances) Re-enable the listeners.
- Configure your system to support the latest Configuration Master. See Overview of Setting Up Configuration Masters, page 8-2.
- Consider saving your configuration. For more information, see Saving and Importing Configuration Settings, page 13-43.

About Reverting to an Earlier Version of AsyncOS

You can revert to an to a previous qualified version of AsyncOS for emergency uses.

Because upgrades cause one-way transformation of key subsystems, the reversion process is complex and requires qualification by Cisco IronPort Quality Assurance teams. You can revert to one of the previous two versions only. The earliest AsyncOS version supported for this functionality is AsyncOS 6.5; prior versions of AsyncOS are not supported.

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.

Important Note About Reversion Impact

Using the `revert` command on a Cisco IronPort appliance is a very destructive action. This command permanently destroys all configuration logs and databases. In addition, it disrupts mail handling until the appliance is reconfigured. Because this command destroys all configuration, it is highly recommended that you have physical local access to the Cisco IronPort appliance when you want to issue the `revert` command.

Warning **You must have a configuration file for the version you want to revert to. Configuration files are *not* backwards-compatible.**

Reverting AsyncOS

Procedure

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:

```
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 IronPort 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 IronPort Spam Quarantine safelist/blocklist database
   to another machine (if applicable)
- waited for the mail queue to empty

Reverting the device causes an immediate reboot to take place.
After rebooting, the appliance reinitializes itself and reboots again to the desired
version.

Do you want to continue? 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 If the appliance will manage Web Security appliances running AsyncOS 7.5 or higher, add at least one of these appliances and wait a few minutes to allow any URL Category updates to be downloaded from the Web Security 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 IronPort 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 IronPort 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, page 13-19.

About URL Category Set Updates for Cisco IronPort Web Usage Controls

- URL Category Set Updates and Centralized Configuration Management, page 8-22
- URL Category Set Updates and Reports, page 5-26

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

Alerts are email notifications containing information about events occurring on the Cisco IronPort appliance. These events can be of varying levels of importance (or severity) from minor to major and pertain generally to a specific component or feature on your appliance. Alerts are generated by the Cisco IronPort appliance. You can specify, at a much more granular level, which alert messages are sent to which users and for which severity of event they are sent. Manage alerts on the Management Appliance > System Administration > Alerts page in the GUI (or via the `alertconfig` command in the CLI).

Overview of Alerts

The following features control the behavior of email notifications:

- **Alerts:** Create alerts to receive email notifications. An alert consists of alert recipients (the email addresses for receiving alerts), and the alert notification (including the severity and alert type).
- **Alert Settings:** Specify global behavior for the alerting feature, including alert sender (FROM:) address, seconds to wait between sending duplicate alerts, and whether to enable AutoSupport (and optionally send weekly AutoSupport reports).

Alerts: Alert Recipients, Alert Classifications, and Severities

Alerts are email messages or notifications containing information about specific functions, such as hardware problems, that are sent to an alert recipient. An alert recipient is an email address to which the alert notifications are sent. The information contained in the notification is determined by the alert classification and severity. You can specify which alert classifications, at which severity, are sent to a particular alert recipient. The alerting engine allows for granular control over the alerts that are sent to recipients. For example, you can configure the system to send only specified types of alerts to a recipient,

such as when the severity level is Critical and the alert type is System. You can also configure general settings (see Configuring Alert Settings, page 13-32). See Alert Listing, page 13-33 for a complete list of alerts.

Alert Classifications

AsyncOS sends the following alert classifications:

- System
- Hardware

Severities

Alerts can be sent for 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 Settings

Alert settings control the general behavior and configuration of alerts, including:

- The RFC 2822 Header From: when sending alerts (enter an address or use the default "alert@<hostname>"). You can also set this via the CLI, using the `alertconfig -> from` command.
- The initial number of seconds to wait before sending a duplicate alert.
- The maximum number of seconds to wait before sending a duplicate alert.
- The status of AutoSupport (enabled or disabled).
- The sending of AutoSupport's weekly status reports to alert recipients set to receive system alerts at the Information level.

Sending Duplicate Alerts

You can specify the initial number of seconds to wait before AsyncOS will send a duplicate alert. If you set this value to 0, duplicate alert summaries are not sent; 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.

Alert Delivery

Because alert messages can be used to inform you of problems within your Cisco IronPort appliance, they are not sent using AsyncOS's normal mail delivery system. Instead, alert messages pass through a separate and parallel email system designed to operate even in the face of significant system failure in AsyncOS.

The alert mail system does not share the same configuration as AsyncOS, which means that alert messages may behave slightly differently from other mail delivery:

- Alert messages are delivered using standard DNS MX and A record lookups.
  - They do not use SMTP routes in AsyncOS versions older then 5.X.
  - They do cache the DNS entries for 30 minutes and the cache is refreshed every 30 minutes, so in case of DNS failure the alerts still go out.
- Alert messages do not pass through the work queue, so they are not scanned for viruses or spam. They are also not subjected to message filters or content filters.
- Alert messages do not pass through the delivery queue, so they will not be affected by bounce profiles or destination control limits.

Cisco IronPort AutoSupport

To allow Cisco IronPort to better support and design future system changes, the Cisco IronPort appliance can be configured to send Cisco IronPort Systems a copy of all alert messages generated by the system. This feature, called 'AutoSupport', is a useful way to allow Cisco IronPort 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 IronPort. This can be disabled if you do not want to send the weekly alert messages internally. To enable or disable this feature, see Configuring Alert Settings, page 13-32.

Alert Messages

Alert messages are standard email messages. Although you can configure the Header From: address, the rest of the message is generated automatically.

Alert From Address

You can configure the Header From: address by clicking the **Edit Settings** button in the GUI or by using the CLI (see the *Cisco IronPort AsyncOS CLI Reference Guide*).

Alert Subject

An alert message's subject has the following format:

```
Subject: [severity]-[hostname]: ([class]) short message
```

Example Alert Message

```
Date: 23 Mar 2007 21:10:19 +0000
To: joe@example.com
From: Cisco IronPort M650 Alert [alert@example.com]
Subject: Critical-example.com: (AntiVirus) update via http://newproxy.example.com failed

The Critical message is:
```

```
update via http://newproxy.example.com failed

Version: 6.0.0-419
Serial Number: XXXXXXXXXXXX-XXXXXXX
Timestamp: Tue May 10 09:39:24 2007

For more information about this error, please see
    http://support.ironport.com
If you need further information, contact your support provider.
```

Managing Alert Recipients

Log in to the GUI and select System Administration > Alerts. (For information about how to access the GUI, see Accessing the Security Management Appliance, page 2-6.)

Figure 13-11 Alerts Page

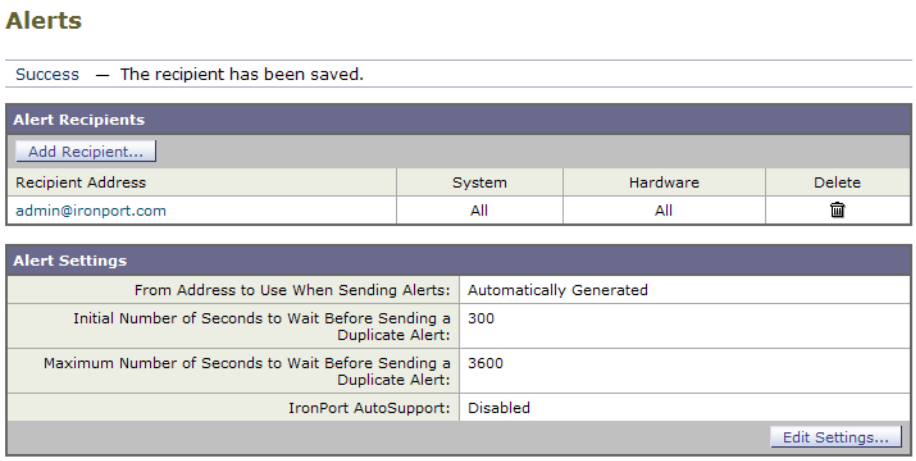

Note If you enabled AutoSupport during system setup, the email address that you specified will receive alerts for all severities and classes by default. You can change the configuration at any time.

The Alerts page lists the existing alert recipients and alert settings.

From the Alerts page, you can:

- Add, configure, or delete alert recipients.
- Modify the alert settings.

Adding New Alert Recipients

Procedure

Step 1 Click **Add Recipient** on the Alerts page.

Step 2 Enter the recipient's email address. You can enter multiple addresses, separated by commas.

Step 3 Select the alert severities that the alert recipient will receive.

Step 4 Click **Submit** to add the alert recipient.

Step 5 Commit your changes.

Configuring Existing Alert Recipients

Procedure

Step 1 Click the alert recipient in the Alert Recipients listing.

Step 2 Make changes to the alert recipient.

Step 3 Submit and commit your changes.

Deleting Alert Recipients

Procedure

Step 1 Click the trash can icon corresponding to the alert recipient in the Alert Recipient listing.

Step 2 Confirm the deletion by clicking **Delete** in the warning dialog that appears.

Step 3 Commit your changes.

Configuring Alert Settings

Alert settings apply to all alerts that the Security Management appliance sends.

Editing Alert Settings

Procedure

Step 1 Click **Edit Settings** on the Alerts page.

Step 2 Enter a Header From: address to use when sending alerts, or select "Automatically generated" ("alert@<hostname>").

Step 3 Select the check box if you want to specify the number of seconds to wait between sending duplicate alerts. For more information, see Sending Duplicate Alerts, page 13-29.

- Specify the initial number of seconds to wait before sending a duplicate alert.
- Specify the maximum number of seconds to wait before sending a duplicate alert.

Step 4 Optionally, enable AutoSupport by selecting the Cisco IronPort AutoSupport option. For more information about AutoSupport, see Cisco IronPort AutoSupport, page 13-30.

If AutoSupport is enabled, the weekly AutoSupport report is sent to alert recipients set to receive system alerts at the Information level. You can disable this via the check box.

Step 5 Submit and commit your changes.

Alert Listing

The following tables list alerts by classification, including the alert name, description, and severity.

Hardware Alerts

Table 13-3 contains a list of the various hardware alerts that AsyncOS can generate, including a description of the alert and the alert severity.

Table 13-3 Listing of Hardware Alerts

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 Alerts

Table 13-4 contains a list of the various system alerts that AsyncOS can generate, including a description of the alert and the alert severity.

Table 13-4 Listing of System Alerts

Alert Name	Description	Severity
COMMON.APP_FAILURE	Sent when there is an unknown application failure.	Critical
COMMON.KEY_EXPIRED_AL ERT	Sent when a feature key has expired.	Warning
COMMON.KEY_EXPIRING_A LERT	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
INTERFACE. FAILOVER.FAILURE_ BACKUP_DETECTED	Sent when a backup NIC pairing interface fails.	Warning
INTERFACE. FAILOVER.FAILURE_ BACKUP_RECOVERED	Sent when a NIC pair failover is recovered.	Information

Table 13-4 Listing of System Alerts (continued)

Alert Name	Description	Severity
INTERFACE.FAILOVER. FAILURE_DETECTED	Sent when a NIC pairing failover is detected due to an interface failure.	Critical
INTERFACE.FAILOVER. FAILURE_DETECTED_NO_ BACKUP	Sent when a NIC pairing failover is detected due to an interface failure, but a backup interface is not available.	Critical
INTERFACE.FAILOVER. FAILURE_RECOVERED	Sent when a NIC pair failover is recovered.	Information
INTERFACE.FAILOVER. MANUAL	Sent when a manual failover to another NIC pair is detected.	Information
COMMON.INVALID_FILTER	Sent when an invalid filter is encountered.	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
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 Cisco IronPort 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

Table 13-4 Listing of System Alerts (continued)

Alert Name	Description	Severity
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_D ISABLED_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
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_FA ILURE_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_F AILED	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

Table 13-4 Listing of System Alerts (continued)

Alert Name	Description	Severity
SYSTEM.SERVICE_ TUNNEL.DISABLED	Sent when a tunnel created for Cisco IronPort Support Services is disabled.	Information
SYSTEM.SERVICE_ TUNNEL.ENABLED	Sent when a tunnel created for Cisco IronPort Support Services is enabled.	Information

Changing Network Settings

This section describes the features used to configure the network operation of the Cisco IronPort 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, page 2-8.

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 `sethostname` command sets the name of the Cisco IronPort 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 Cisco IronPort 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 13-5 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 IronPort 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, page 13-37.

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 `dnsflush` command (for more information about the `dnsflush` command, see the *Cisco IronPort AsyncOS CLI Reference Guide*), clears all information in the DNS cache. You may choose to use this feature when changes have been made to your local DNS system. The command takes place immediately and may cause a temporary performance degradation while the cache is repopulated.

Configuring DNS Settings via the Graphical User Interface

Procedure

Step 1 On the **Management Appliance > Network > DNS** page, click the **Edit Settings** button.

Step 2 Select whether to use the Internet's root DNS servers or your own internal DNS server(s), and specify authoritative DNS servers.

Step 3 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, page 13-37.

Step 4 Choose an interface for DNS traffic.

Step 5 Enter the number of seconds to wait before canceling a reverse DNS lookup.

Step 6 Optionally, clear the DNS cache by clicking **Clear Cache.**

Step 7 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.

Adding Static Routes

Procedure

Step 1 On the Management Appliance > Network > Routing page, click **Add Route** in the route listing. The Enter a name for the route.

Step 2 Enter the destination IP address.

Step 3 Enter the gateway IP address.

Step 4 Submit and commit your changes.

Deleting Static Routes

Procedure

Step 1 Click the trash can icon corresponding to the static route name in the Static Routes listing.

Step 2 Confirm the deletion by clicking **Delete** in the warning dialog that appears.

Step 3 Commit your changes.

Editing Static Routes

Procedure

Step 1 Click the name of the route in the Static Route listing.

Step 2 Make changes to the route.

Step 3 Submit and commit your changes.

Modifying the Default Gateway (GUI)

Procedure

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), page 13-40) or via the `setgateway` command in the CLI.

Changing the Admin User's Password

The password for the admin user can be changed via the GUI or the CLI.

To change the password via the GUI, use the Management Appliance > System Administration > Users page. For more information, see Setting and Changing Passwords, page 12-12.

To change the password for the `admin` user in the CLI, use the `password` command. Passwords must be six characters or longer. The `password` command requires you to enter the old password for security.

Note Changes to the password take effect immediately and do not require you to carry out the **commit** command.

Configuring the System Time

You can set the system time on the Cisco IronPort appliance and specify the time zone. Use the Management Appliance > System Administration > Time Zone page and the Management Appliance > System Administration > Time Settings page in the GUI. Alternatively, use the `ntpconfig`, `settime`, and `settz` commands in the CLI.

Time Zone Page

The Time Zone page (available on the System Administration menu in the GUI) displays the time zone for the Cisco IronPort appliance. You can select a specific time zone or GMT offset.

Selecting a Time Zone

To set the time zone for the Cisco IronPort appliance:

Procedure

Step 1 Choose **Management Appliance > System Administration > Time Zone**.

Step 2 Click **Edit Settings**.

Step 3 Select a region, country, and time zone.

Step 4 Submit and commit your changes.

Selecting a GMT Offset

To set the GMT offset for the Cisco IronPort appliance:

Procedure

Step 1 Choose **Management Appliance > System Administration > Time Zone**.

Step 2 Click **Edit Settings**.

Step 3 Select GMT Offset from the list of regions. The Time Zone Setting page is updated to include GMT offsets in the Time Zone field.

Figure 13-12 Setting a GMT Offset

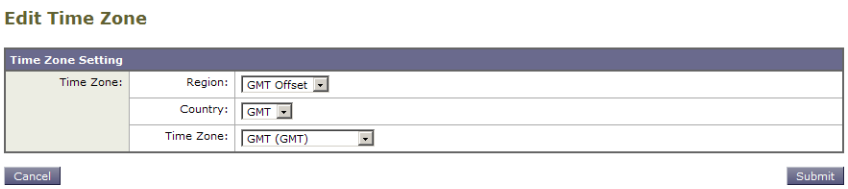

Step 4 Select an offset in the Time Zone field. The offset refers to the number of hours that you add or subtract to or from Greenwich Mean Time (GMT) — the local time at the prime meridian. Hours preceded by a minus sign ("-") are west of the prime meridian. A plus sign ("+") indicates locations east of the prime meridian.

Step 5 Submit and commit your changes.

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 on how the Security Management appliance gathers information, see "How the Security Appliance Gathers Data for Reports" section on page 3-2.

Updating Time Zone Files

Whenever there is a change in the time zone rules for any country, you may need to update the Time Zone files on the appliance.

Automatically Updating Time Zone Files

Procedure

Step 1 Choose **Management Appliance > System Administration > Update Settings**.

Step 2 Select the **Enable automatic updates for Time zone rules** check box.

Step 3 Enter an interval. Click the **?** help on the page for important information.

Step 4 If you have not yet done so, configure the other settings on this page. See Configuring Upgrade and Service Update Settings, page 13-19.

Manually Updating Time Zone Files

Procedure

Step 1 Choose **Management Appliance > System Administration > Time Settings**.

Step 2 Look at the **Time Zone File Updates** section.

Step 3 If there is an available time zone file update, click **Update Now**.

Editing Time Settings

Set the system time on the Security Management appliance.

Setting the System Time Using a Network Time Protocol (NTP) Server

Using an NTP server synchronizes the Security Management appliance system clock with other computers on your network.

Procedure

Step 1 Click **Edit Settings** on the Management Appliance > System Administration > Time Settings page.

Step 2 Select **Use Network Time Protocol**.

Step 3 Enter an NTP server address

Step 4 (Optional) Click **Add Row** to add additional NTP servers.

Step 5 To delete an NTP server from the list, click the trash can icon for that server.

Step 6 Select an interface for NTP queries. This is the IP address from which NTP queries should originate.

Step 7 Submit and commit your changes.

Setting System Time Manually

Procedure

Step 1 Click **Edit Settings** on the Management Appliance > System Administration > Time Settings page.

Step 2 Select **Set Time Manually**.

Step 3 Enter the date in MM/DD/YYYY format, or click the calendar icon to select the date.

Step 4 Enter the local time in HH:MM:SS format.

Step 5 Submit and commit your changes.

Saving and Importing Configuration Settings

Note The configuration file described in this section is used to configure Security Management appliances. The configuration files and Configuration Master described in Chapter 8, "Managing Web Security Appliances" are used to configure Web Security 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.)

Managing Multiple Appliances with XML Configuration Files

Warning **If you want to import a configuration file from one Security Management appliance into another Security Management appliance:**

Everything in the original configuration, including the IP Address, is included in the configuration file. Either edit the configuration file to change the IP address, or be sure the original Security Management appliance is offline.

Also, be aware that the SSH authentication connection will terminate. When this happens, you will need to re-establish the connection for all connected Web Security appliances and Email Security appliances.

- You can download an existing configuration file from one Cisco IronPort appliance, make changes to it, and upload it to a different appliance. This lets you manage an installation of multiple Cisco IronPort appliances more easily. However, you cannot load configuration files from Email Security appliances onto a Security Management appliance.
- You can divide an existing configuration file downloaded from one appliance into multiple subsections. You can modify those sections that are common among all appliances (in a multiple appliance environment) and load them onto other appliances as the subsections are updated.

  For example, you could use an appliance in a test environment for testing the Global Unsubscribe command. When you feel that you have configured the Global Unsubscribe list appropriately, you could then load the Global Unsubscribe configuration section from the test appliance to all of your production appliances.

Managing Configuration Files

To manage configuration files on your appliance, choose **Management Appliance > System Administration > Configuration File**.

The Configuration File page contains the following sections:

- **Current Configuration:** used to save and export the current configuration file
- **Load Configuration:** used to load a complete or partial configuration file
- **End-User Safelist/Blocklist Database (Cisco IronPort Spam Quarantine):** used to manage the safelist/blocklist database
- **Reset Configuration:** used to reset the current configuration back to the factory defaults (you should save your configuration prior to resetting it)

Saving and Exporting the Current Configuration File

Using the Current Configuration section of the 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

Use the Load Configuration section of the Management Appliance > System Administration > Configuration File page to load new configuration information into the Cisco IronPort appliance. You can load information using one of three methods:

- Placing information in the `configuration` directory and uploading it
- Uploading the configuration file directly from your local machine
- Pasting configuration information directly into the GUI

Configuration files with masked passwords cannot be loaded.

Figure 13-13 Loading a Configuration File

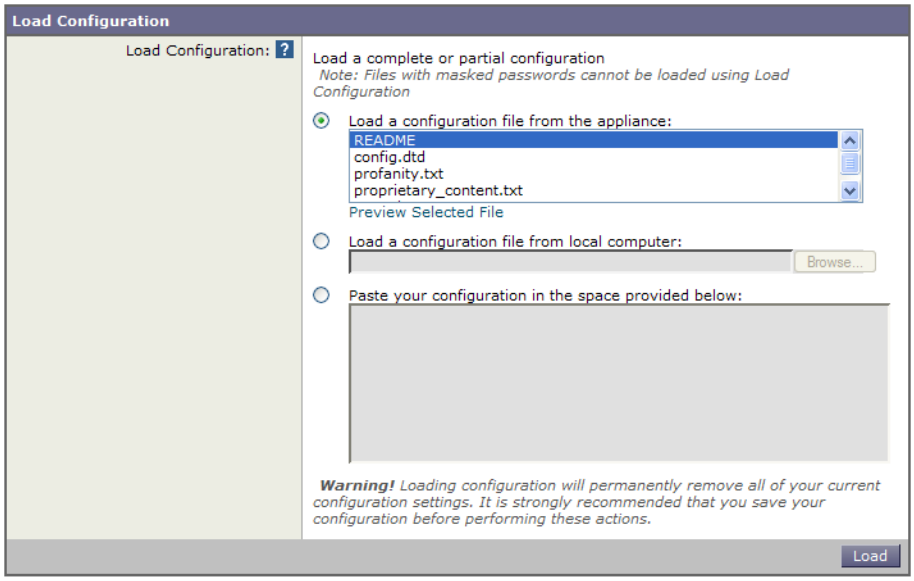

Regardless of the method, you must include the following tags at the top of your configuration:

```
<?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 IronPort appliance. The DTD file is named `config.dtd`. If validation errors are reported at the command line when you use the `loadconfig` command, the changes are not loaded. You can download the DTD to validate configuration files outside of the appliance before uploading them.

In either method, you can import an entire configuration file (the information defined between the highest level tags: `<config></config>`), or a *complete* and *unique* 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.

Warning **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>
```

Warning **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"?>
```

Currently, only configuration files with this encoding can be loaded.

Resetting the Current Configuration

Resetting the current configuration causes your Cisco IronPort appliance to revert back to the original factory defaults. Save your configuration prior to resetting it. Resetting the configuration via the **Reset** button in the GUI is not supported in a clustering environment.

Figure 13-14 Resetting the Configuration File

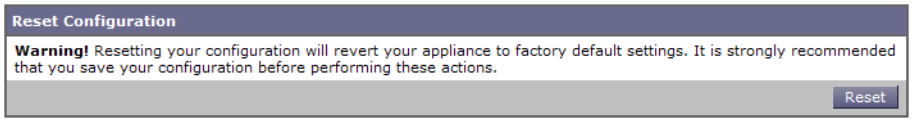

See Resetting to Factory Defaults, page 13-5.

CLI Commands for Configuration Files

The following commands enable you to manipulate the configuration files:

- `showconfig`
- `mailconfig`
- `saveconfig`
- `loadconfig`
- `resetconfig` (see Resetting to Factory Defaults, page 13-5)

- publishconfig
- backupconfig

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, page 13-47.

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: Cisco IronPort model number Messaging Gateway Appliance(tm)
  Model Number: model number
  Version: version of AsyncOS installed
  Serial Number: serial number
  Current Time: current time and date
[The remainder of the configuration file is printed to the screen.]
```

Use the mailconfig command to email the current configuration to a user. A configuration file in XML format named config.xml will be attached to the message.

```
mail3.example.com> mailconfig

Please enter the email address to which you want to send
the configuration file.
[]> administrator@example.com

Do you want to include passwords? Please be aware that a configuration without passwords
will fail when reloaded with loadconfig.  [N]> y

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

The saveconfig command on the Security Management appliance stores and saves all of the configuration master files (ESA and WSA) 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 Cisco IronPort 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, page 13-45 for more information.

The publishconfig Command

Use the `publishconfig` command to publish changes a configuration master. The syntax is as follows:

publishconfig *config_master* [*job_name*] [*host_list* | *host_ip*

where *config_master* is a supported Configuration Master, as listed in the Compatibility Matrix in the Release Notes for this release at http://www.cisco.com/en/US/products/ps10155/prod_release_notes_list.html. This keyword is required. The keyword *job_name* is optional and will be generated if it is not specified.

The keyword *host_list* is a list of host names or IP addresses for WSA appliances to be published, and will be published to all hosts assigned to the configuration master if not specified. The optional *host_ip* can be multiple host IP addresses, each separated by a comma.

To verify that the `publishconfig` command was successful, check the `smad_logs` file. You can also verify that the publish history was successful from the Security Management appliance GUI by choosing **Web > Utilities > Web Appliance Status**. From this page choose the web appliance that you want the publish history details. Additionally, you can go the Publish History page: **Web > Utilities > Publish > Publish History**.

The backupconfig Command

Use the `backupconfig` command to copy the active dataset from the 'source' appliance to a 'target' Security Management appliance with minimum disruption on the originating 'source' Security Management appliance.

For more information on this command, the usage, and backing up your dataset, see "Backing Up Your Security Management Appliance" section on page 13-6.

Uploading Configuration Changes Using the CLI

Procedure

Step 1 Outside of the CLI, ensure that you are able to access the `configuration` directory of the appliance. See Appendix A, "IP Interfaces and Accessing the Appliance" 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:

```
mail3.example.com> loadconfig

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, page 2-8.) Finally, the changes are committed.

```
mail3.example.com> loadconfig

1. Paste via CLI
2. Load from file
[1]> 1

Paste the configuration file now. Press CTRL-D on a blank line when done.

[The configuration file is pasted until the end tag </config>. Control-D is entered on a
separate line.]

Values have been loaded.
Be sure to run "commit" to make these settings active.

mail3.example.com> commit

Please enter some comments describing your changes:
[]> pasted new configuration file and changed default settings
```

Managing Disk Usage

View the amount of disk space allocated to each of the Security Management appliance's monitoring services at **Management Appliance > System Administration > Disk Management**.

Disk Space Maximums and Allocations

You can allocate available disk space among the features that your organization uses, up to the maximum available.

Table 13-6 Maximum Disk Space Available, in GB

	Hardware Platform					
	M160	M170	M660	M670	M1060	M1070
Total available for all features including Spam Quarantine	165	165	681	681	1039	1407
Spam Quarantine Maximum	70	70	150	150	265	265

Table 13-7 Default Disk Space Allocations by Feature, in Percent

Feature	Disk Space Allocated by Default
Centralized Reporting (Email and Web)	10 %
Email Tracking	30 %
Web Tracking	30 %
Spam Quarantine	30 %

Note

- Unlike reporting (which is just counters) and tracking (which stores limited amounts of header information), the spam quarantine stores the entire message bodies of quarantined messages, and therefore uses significantly more space per message than the other features. Because of this significant space usage, the spam quarantine disk quota has stricter limitations than just the available disk space, in order to avoid locking up the appliance.
- Centralized Reporting Disk Space on Security Management appliances is used for both Email and Web data. If you enable either Centralized Email Reporting or Centralized Web Reporting, all of the space is dedicated to the enabled feature. If you enable both, Email and Web reporting data share the space and space is allocated on a first-come basis.
- If you enable centralized web reporting but there is no disk space allocated for reporting, then centralized web reporting will not work until disk space is allocated.
- If you reduce the existing allocation, 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.

Reallocating Disk Space Quotas

If the appliance frequently runs out of disk space for a particular feature and has excess space for other features, you can reallocate disk quotas to relieve the problem. If you require more space for all features, consider upgrading your hardware.

Before You Begin

- Changing disk allocations may impact existing data or feature availability. See information at Disk Space Maximums and Allocations, page 13-50.

- You can temporarily create space in a quarantine by manually releasing or deleting messages from the quarantine.

Procedure

Step 1 On the Security Management appliance, choose **Management Appliance > System Administration > Disk Management**

Step 2 Click **Edit Disk Quotas.**

Step 3 On the **Edit Disk Quotas** page, enter the amount of disk space (in gigabytes) allocated to each service.

For all services except Spam Quarantine, you can enter a value between 0 and the total amount of disk space.

Step 4 Click **Submit.**

Step 5 In the confirmation dialog box, click **Set New Quotas.**

Step 6 Click **Commit** to commit your changes.

Customizing Your View

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 and to PDF reports)
- 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 and Web 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 and commit your changes when done.

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.

CHAPTER 14

Logging

- Logging Overview, page 14-1
- Log Types, page 14-4
- Log Subscriptions, page 14-22

Logging Overview

Log files record regular operations, as well as exceptions, for activity on the system. Use the logs for monitoring the Cisco IronPort 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 Security Management appliance.

Note The Security Management appliance pulls information for all reporting and tracking with the exception of Cisco IronPort Spam Quarantine (ISQ) data. The ISQ data is pushed from the ESA.

Log Retrieval

Log files can be retrieved with the file transfer protocols described in Table 14-1. 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 14-1 Log Transfer Protocols

FTP Poll	With this type of file transfer, a remote FTP client accesses the Cisco IronPort 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 IronPort 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 IronPort 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 IronPort 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 Cisco IronPort 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 14-2 Preconfigured Log Subscriptions

Log Name	Log Type	Retrieval Method
`cli_logs`	CLI Audit Logs	FTP Poll
`euq_logs`	Cisco IronPort Spam Quarantine Logs	FTP Poll
`euqgui_logs`	Cisco IronPort Spam Quarantine GUI Logs	FTP Poll
`gui_logs`	HTTP Logs	FTP Poll
`mail_logs`	Cisco IronPort 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

Table 14-2 Preconfigured Log Subscriptions (continued)

Log Name	Log Type	Retrieval Method
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, page 14-23.

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, page 14-22.

Log Types

- Summary of Log Types, page 14-4
- Using Configuration History Logs, page 14-7
- Using CLI Audit Logs, page 14-8
- Using FTP Server Logs, page 14-9
- Using HTTP Logs, page 14-9
- Using Cisco IronPort Spam Quarantine Logs, page 14-10
- Using Cisco IronPort Spam Quarantine GUI Logs, page 14-10
- Using Cisco IronPort Text Mail Logs, page 14-11
- Using NTP Logs, page 14-16
- Using Reporting Logs, page 14-16
- Using Reporting Query Logs, page 14-17
- Using Safelist/Blocklist Logs, page 14-17
- Using SMA Logs, page 14-18
- Using Status Logs, page 14-19
- Using System Logs, page 14-21
- Understanding Tracking Logs, page 14-21

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, page 14-22 for more information.

AsyncOS generates the following log types:

Table 14-3 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.
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, page 6-1.
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, page 13-41.
Reporting Logs	Reporting logs record actions associated with the processes of the centralized reporting service.

Table 14-3 Log Types (continued)

Log Type	Description
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 Cisco IronPort 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.
Spam Quarantine GUI Logs	Cisco IronPort Spam Quarantine GUI logs record actions associated with the Cisco IronPort Spam Quarantine GUI, such as quarantine configuration through the GUI, end user authentication, and end user actions (for example, releasing email).
Spam Quarantine Logs	Cisco IronPort Spam Quarantine logs record actions associated with the Cisco IronPort 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.

Log Type Comparison

Table 14-4 summarizes the characteristics of each log type.

Table 14-4 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	•		•								
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	•		•			•					
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.

Configuration History Log Example

In this example, the configuration history log shows that the user (admin) added a guest user to the table that defines which local users are allowed to log in to the system.

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<!DOCTYPE config SYSTEM "config.dtd">
<!--
  XML generated by configuration change.
  Change comment: added guest user
  User: admin
  Configuration are described as:
    This table defines which local users are allowed to log into the system.
  Product: Cisco IronPort M160 Messaging Gateway(tm) Appliance
  Model Number: M160
  Version: 6.7.0-231
  Serial Number: 000000000ABC-D000000
  Number of CPUs: 1
  Memory (GB): 4
  Current Time: Thu Mar 26 05:34:36 2009
  Feature "Cisco IronPort Centralized Configuration Manager": Quantity = 10, Time
Remaining = "25 days"
  Feature "Centralized Reporting": Quantity = 10, Time Remaining = "9 days"
  Feature "Centralized Tracking": Quantity = 10, Time Remaining = "30 days"
  Feature "Centralized Spam Quarantine": Quantity = 10, Time Remaining = "30 days"
  Feature "Receiving": Quantity = 1, Time Remaining = "Perpetual"
-->
<config>
```

Using CLI Audit Logs

Table 14-5 describes the statistics recorded in CLI audit logs.

Table 14-5 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.

CLI Audit Log Example

In this example, the CLI audit log shows that, for PID 16434, the following CLI commands were entered: `who`, `textconfig`.

```
Thu Sep  9 14:35:55 2004 Info: PID 16434: User admin entered 'who'; prompt was
'\nmail3.example.com> '
Thu Sep  9 14:37:12 2004 Info: PID 16434: User admin entered 'textconfig'; prompt was
'\nUsername  Login Time  Idle Time  Remote Host  What\n
========  ==========  =========  ===========  ====\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

Table 14-6 describes the statistics recorded in FTP server logs.

Table 14-6 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).

FTP Server Log Example

In this example, the FTP server log records a connection (ID:1). The IP address of the incoming connection is shown, as well as the activity (uploading and downloading files) and the logout.

```
Wed Sep  8 18:03:06 2004 Info: Begin Logfile
Wed Sep  8 18:03:06 2004 Info: Version: 4.0.0-206 SN: 00065BF3BA6D-9WFWC21
Wed Sep  8 18:03:06 2004 Info: Time offset from UTC: 0 seconds
Wed Sep  8 18:03:06 2004 Info: System is coming up
Fri Sep 10 08:07:32 2004 Info: Time offset from UTC: -25200 seconds
Fri Sep 10 08:07:32 2004 Info: ID:1 Connection from 10.1.3.14 on 172.19.0.86
Fri Sep 10 08:07:38 2004 Info: ID:1 User admin login SUCCESS
Fri Sep 10 08:08:46 2004 Info: ID:1 Upload wording.txt 20 bytes
Fri Sep 10 08:08:57 2004 Info: ID:1 Download words.txt 1191 bytes
Fri Sep 10 08:09:06 2004 Info: ID:1 User admin logout
```

Using HTTP Logs

Table 14-7 describes the statistics recorded in HTTP logs

Table 14-7 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.

HTTP Log 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 Cisco IronPort Spam Quarantine Logs

Table 14-8 describes the statistics recorded in Cisco IronPort Spam Quarantine logs.

Table 14-8 Cisco IronPort 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).

Cisco IronPort Spam Quarantine Log 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 Cisco IronPort Spam Quarantine GUI Logs

Table 14-9 shows the statistics recorded in Cisco IronPort Spam Quarantine GUI logs.

Table 14-9 Cisco IronPort 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.

Cisco IronPort Spam Quarantine GUI Log Example

In this example, the log shows a successful authentication, login, and logout:

Table 14-10 Cisco IronPort 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 Cisco IronPort 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, page 6-1.

Table 14-11 shows the information displayed in text mail logs.

Table 14-11 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 Cisco IronPort Spam Quarantine. It is used to track messages as they are sent to and from the Cisco IronPort Spam Quarantine.
MID	Message ID: Use this to track messages as they flow through the logs.
RID	Recipient ID. Each message recipient is assigned an ID.
New	New connection initiated.
Start	New message started.

Sample

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 14-12 Text Mail Log Detail

```
1  Mon Apr 17 19:56:22 2003 Info: New SMTP ICID 5 interface Management
   (10.1.1.1) address 10.1.1.209 reverse dns host remotehost.com
   verified yes
2  Mon Apr 17 19:57:20 2003 Info: Start MID 6 ICID 5
3  Mon Apr 17 19:57:20 2003 Info: MID 6 ICID 5 From:
   <sender@remotehost.com>
4  Mon Apr 17 19:58:06 2003 Info: MID 6 ICID 5 RID 0 To:
   <mary@yourdomain.com>
5  Mon Apr 17 19:59:52 2003 Info: MID 6 ready 100 bytes from
   <sender@remotehost.com>
6  Mon Apr 17 19:59:59 2003 Info: ICID 5 close
7  Mon Mar 31 20:10:58 2003 Info: New SMTP DCID 8 interface
   192.168.42.42 address 10.5.3.25
8  Mon Mar 31 20:10:58 2003 Info: Delivery start DCID 8 MID 6 to RID
   [0]
9  Mon Mar 31 20:10:58 2003 Info: Message done DCID 8 MID 6 to RID [0]
10 Mon Mar 31 20:11:03 2003 Info: DCID 8 close
```

Use Table 14-13 as a guide to reading the previous log file.

Table 14-13 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 Cisco IronPort 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 Cisco IronPort appliance. Upon delivery, the destination host returns a 5*XX* error, which indicates that the message cannot be delivered to either recipient. The Cisco IronPort 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 Cisco IronPort 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 Cisco IronPort 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 Cisco IronPort 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 Cisco
IronPort Spam Quarantine
```

```
Wed Feb 14 12:11:45 2007 Info: EUQ: Quarantined MID 2317877
Wed Feb 14 12:11:45 2007 Info: RPC Message done RCID 756814 MID 2317877
Wed Feb 14 12:11:45 2007 Info: Message finished MID 2317877 done
```

Using NTP Logs

Table 14-14 shows the statistics recorded in NTP logs.

Table 14-14 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.

NTP Log Example

In this example, the NTP log shows the appliance polling the NTP host twice.

```
Thu Sep  9 07:36:39 2004 Info: sntp query host 10.1.1.23 delay 653 offset -652
Thu Sep  9 07:36:39 2004 Info: adjust: time_const: 8 offset: -652us next_poll: 4096
Thu Sep  9 08:44:59 2004 Info: sntp query host 10.1.1.23 delay 642 offset -1152
Thu Sep  9 08:44:59 2004 Info: adjust: time_const: 8 offset: -1152us next_poll: 4096
```

Using Reporting Logs

Table 14-15 shows the statistics recorded in reporting logs.

Table 14-15 Reporting Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken, including user authentication, and so forth.

Reporting Log Example

In this example, the Reporting log shows the appliance set at the information log level.

```
Wed Oct  3 13:39:53 2007 Info: Period minute using 0 (KB)
Wed Oct  3 13:39:53 2007 Info: Period month using 1328 (KB)
Wed Oct  3 13:40:02 2007 Info: Update 2 registered appliance at 2007-10-03-13-40
Wed Oct  3 13:40:53 2007 Info: Pages found in cache: 1304596 (99%). Not found: 1692
Wed Oct  3 13:40:53 2007 Info: Period hour using 36800 (KB)
Wed Oct  3 13:40:53 2007 Info: Period day using 2768 (KB)
Wed Oct  3 13:40:53 2007 Info: Period minute using 0 (KB)
Wed Oct  3 13:40:53 2007 Info: Period month using 1328 (KB)
Wed Oct  3 13:40:53 2007 Info: HELPER checkpointed in 0.00580507753533 seconds
Wed Oct  3 13:41:02 2007 Info: Update 2 registered appliance at 2007-10-03-13-41
Wed Oct  3 13:41:53 2007 Info: Pages found in cache: 1304704 (99%). Not found: 1692
Wed Oct  3 13:41:53 2007 Info: Period hour using 36800 (KB)
Wed Oct  3 13:41:53 2007 Info: Period day using 2768 (KB)
Wed Oct  3 13:41:53 2007 Info: Period minute using 0 (KB)
```

```
Wed Oct  3 13:41:53 2007 Info: Period month using 1328 (KB)
Wed Oct  3 13:42:03 2007 Info: Update 2 registered appliance at 2007-10-03-13-42
```

Using Reporting Query Logs

Table 14-16 shows the statistics recorded in reporting query logs.

Table 14-16 Reporting Query Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken, including user authentication, and so forth.

Reporting Query Log Example

In this example, the reporting query log shows the appliance running a daily outgoing email traffic query for the period from August 29 to October 10, 2007.

```
Tue Oct  2 11:30:02 2007 Info: Query: Closing interval handle 811804479.
Tue Oct  2 11:30:02 2007 Info: Query: Closing interval handle 811804480.
Tue Oct  2 11:30:02 2007 Info: Query: Closing query handle 302610228.
Tue Oct  2 11:30:02 2007 Info: Query: Merge query with handle 302610229 for
['MAIL_OUTGOING_TRAFFIC_SUMMARY.
DETECTED_SPAM', 'MAIL_OUTGOING_TRAFFIC_SUMMARY.DETECTED_VIRUS',
'MAIL_OUTGOING_TRAFFIC_SUMMARY.THREAT_CONTEN
T_FILTER', 'MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_CLEAN_RECIPIENTS',
'MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_RECI
PIENTS_PROCESSED'] for rollup period "day" with interval range 2007-08-29 to 2007-10-01
with key constraints
 None sorting on ['MAIL_OUTGOING_TRAFFIC_SUMMARY.DETECTED_SPAM'] returning results from 0
to 2 sort_ascendin
g=False.
Tue Oct  2 11:30:02 2007 Info: Query: Closing query handle 302610229.
Tue Oct  2 11:30:02 2007 Info: Query: Merge query with handle 302610230 for
['MAIL_OUTGOING_TRAFFIC_SUMMARY.
TOTAL_HARD_BOUNCES', 'MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_RECIPIENTS_DELIVERED',
'MAIL_OUTGOING_TRAFFIC_SUMM
ARY.TOTAL_RECIPIENTS'] for rollup period "day" with interval range 2007-08-29 to
2007-10-01 with key constra
ints None sorting on ['MAIL_OUTGOING_TRAFFIC_SUMMARY.TOTAL_HARD_BOUNCES'] returning
results from 0 to 2 sort
_ascending=False.
Tue Oct  2 11:30:02 2007 Info: Query: Closing query handle 302610230.
```

Using Safelist/Blocklist Logs

Table 14-17 shows the statistics recorded in safelist/blocklist logs.

Table 14-17 Safelist/Blocklist Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The message consists of actions taken, including user authentication, and so forth.

Safelist/Blocklist Log Example

In this example, the safelist/blocklist log shows the appliance creating database snapshots every two hours. It also shows when senders were added to the database.

```
Fri Sep 28 14:22:33 2007 Info: Begin Logfile Fri Sep 28 14:22:33 2007 Info: Version:
6.0.0-425 SN: XXXXXXXXXXX-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

Table 14-18 shows the statistics recorded in SMA logs.

Table 14-18 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.

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.

Reading Status Logs

Table 14-19 shows the status log labels and the matching system statistics.

Table 14-19 Status Log Statistics

Statistic	Description
CPULd	CPU utilization.
DskIO	Disk I/O utilization.
RAMUtil	RAM utilization.
QKUsd	Queue kilobytes used.
QKFre	Queue kilobytes free.
CrtMID	Message ID (MID).
CrtICID	Injection connection ID (ICID).
CRTDCID	Delivery connection ID (DCID).
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.
UnatmptRcp	Unattempted recipients.
AtmptRcp	Attempted recipients.

Table 14-19 Status Log Statistics (continued)

Statistic	Description
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.
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 Cisco IronPort Spam Quarantine.
EuqRls	Number of messages in the Cisco IronPort Spam Quarantine release queue.

Status Log Example

```
Fri Feb 24 15:14:39 2006 Info: Status: CPULd 0 DskIO 0 RAMUtil 2 QKUsd 0 QKFre 8388608
CrtMID 19036 CrtICID 35284 CrtDCID 4861 InjMsg 13889 InjRcp 14230 GenBncRcp 12 RejRcp 6318
DrpMsg 7437 SftBncEvnt 1816 CmpRcp 6813 HrdBncRcp 18 DnsHrdBnc 2 5XXHrdBnc 15 FltrHrdBnc 0
ExpHrdBnc 1 OtrHrdBnc 0 DlvRcp 6793 DelRcp 2 GlbUnsbHt 0 ActvRcp 0 UnatmptRcp 0 AtmptRcp 0
CrtCncIn 0 CrtCncOut 0 DnsReq 143736 NetReq 224227 CchHit 469058 CchMis 504791 CchEct
15395 CchExp 55085 CPUTTm 228 CPUETm 181380 MaxIO 350 RAMUsd 21528056 MMLen 0 DstInMem 4
ResCon 0 WorkQ 0 QuarMsgs 0 QuarQKUsd 0 LogUsd 3 AVLd 0 BMLd 0 CASELd 3 TotalLd 3 LogAvail
17G EuQ 0 EuqRls 0
```

Using System Logs

Table 14-20 shows the statistics recorded in system logs.

Table 14-20 System Log Statistics

Statistic	Description
Timestamp	Time that the bytes were transmitted.
Message	The logged event.

System Log Example

In this example, the system log shows some commit entries, including the name of the user issuing the commit and the comment entered.

```
Wed Sep  8 18:02:45 2004 Info: Version: 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 IronPort. The conversion tools are located at the following URL:

`http://tinyurl.com/3c518r`.

Log Subscriptions

- Configuring Log Subscriptions, page 14-22
- Creating a Log Subscription in the GUI, page 14-23
- Configuring Global Settings for Logging, page 14-24
- Rolling Over Log Subscriptions, page 14-26
- Configuring Host Keys, page 14-28

Configuring Log Subscriptions

Log subscriptions create the individual log files that are stored on a Cisco IronPort 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 14-21 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, page 14-4.
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 Cisco IronPort 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, page 14-2.
Log Level	Level of detail for each log subscription.
Retrieval Method	Method used to transfer the log files from the Cisco IronPort 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, page 14-4. 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 14-22 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.
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

Procedure

Step 1 On the Management Appliance > System Administration > Log Subscriptions page, click **Add Log Subscription**.

Step 2 Select a log type and enter the log name (for the log directory), as well as the name for the log file itself.

Step 3 If applicable, specify the maximum file size.

Step 4 If applicable, specify days, times of day, or time intervals to roll over the logs. For more information, see Log Rollover and Transfer Schedule, page 14-2.

Step 5 If applicable, specify the log level.

Step 6 (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 7 Configure the log retrieval method.

Step 8 Submit and commit your changes.

Editing Log Subscriptions

Procedure

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 IronPort 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 IronPort 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 IronPort 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 Cisco IronPort 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 14-23 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

Procedure

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, page 14-24.

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)

Rolling Over Logs in Log Subscriptions

See Log Rollover and Transfer Schedule, page 14-2.

Rolling Over Logs Immediately Using the GUI

Procedure

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.

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.

Figure 14-1 Viewing Log Files in the GUI

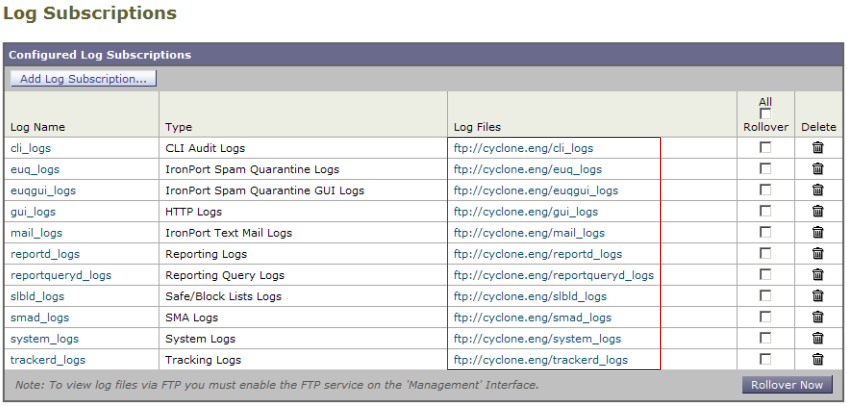

Viewing the Most Recent Entries in Logs (tail Command)

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 Cisco IronPort 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: "Cisco IronPort Spam Quarantine Logs" Retrieval: FTP Poll
3. "euqgui_logs" Type: "Cisco IronPort Spam Quarantine GUI Logs" Retrieval: FTP Poll
4. "gui_logs" Type: "HTTP Logs" Retrieval: FTP Poll
5. "mail_logs" Type: "Cisco IronPort 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 IronPort 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 *Cisco IronPort AsyncOS for Email Security Daily Management Guide*.

The `hostkeyconfig` subcommand performs the following functions:

Table 14-24 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.

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 15

Troubleshooting

- Accessing Customer Support from the Security Management Appliance, page 15-1
- Packet Capture, page 15-3

Accessing Customer Support from the Security Management Appliance

The following commands and features are useful when you need to contact Customer Support or activate features on the Security Management appliance:

- Technical Support, page 15-1
- Packet Capture, page 15-3

Technical Support

Use the **Help and Support** menu in the upper-right corner of the GUI to access the features related to Cisco IronPort Customer Support.

The technical support features include two pages: the **Open a Support Case** page and the **Remote Access** page.

Creating a Support Request

You can use the Help and **Support > Open a Support Case** page or the **supportrequest** command to email the configuration of your appliance to Customer Support or to other users, and enter comments describing the issue for which you need support. See the *Cisco IronPort AsyncOS CLI Reference Guide* for more information about the **supportrequest** command. This command requires that the appliance is able to send mail to the Internet.

Procedure

Step 1 Choose **Help and Support > Open a Support Case**.

Step 2 Enter your contact information (name, email address, and phone numbers).

Step 3 Enter a description of the issue.

Step 4 Optionally, enter additional recipients' email addresses in the **Other recipients** field.

By default, the support request (including the configuration file) is sent to Cisco IronPort Customer Support when you select the check box at the top of the form. You can also mail the configuration file to other email addresses. Separate multiple addresses with commas.

Step 5 If you already have a customer support ticket for this issue, enter the ticket number at the bottom of the page.

Step 6 Click **Send.**

A trouble ticket is automatically created. For additional information, see Customer Support, page 1-4.

About Remote Access for Customer Support

Enabling Remote Access activates a special account used by Customer Support for debugging and general access to the system. This is used by Cisco IronPort Customer Support for tasks such as helping you configure your system, understanding your configurations, and investigating problem reports. You can also use the **techsupport** command in the CLI.

When enabling the use of the Secure Tunnel, the appliance creates an SSH tunnel over the specified port to the server **upgrades.cisco.com**. By default this connection is over port 25, which works in most environments because the system also requires general access over that port in order to send email messages. After a connection is made to **upgrades.cisco.com**, Customer Support is able to use the SSH tunnel to obtain access to the appliance. As long as the connection over port 25 is allowed, this bypasses most firewall restrictions. You can also use the **techsupport tunnel** command in the CLI.

In both the Remote Access mode and Tunnel mode, a password is required. This is *not* the password that will be used to access the system. After this password and the system serial number are provided to your customer support representative, a password used to access the appliance is generated.

After the technical support tunnel is enabled, it remains connected to **upgrades.cisco.com** for seven days. At the end of the seven days, established connections are not disconnected but are unable to reattach to the tunnel once disconnected. The timeout set on the SSH tunnel connection does not apply to the Remote Access account; it remains active until specifically deactivated.

Enabling Remote Access for Customer Support

Procedure

Step 1 On the right hand side of the Security Management appliance GUI, choose **Help and Support> Remote Access**.

Step 2 Select **Edit Remote Access Settings**.

Step 3 Check the **Allow remote access to this appliance** check box.

Step 4 Enter your customer support password.

Step 5 Leave the **Secure Tunnel** check box selected and leave the port number 25, unless a customer support engineer instructs you to change these options.

Step 6 Submit and commit your changes.

Packet Capture

Sometimes when you contact Cisco IronPort Customer Support with a Security Management appliance issue, you may be asked to provide insight into the network activity going into and out of the Security Management appliance. The Security Management appliance provides the ability to intercept and display TCP/IP and other packets being transmitted or received over the network to which the appliance is attached.

You might want to run a packet capture to debug the network setup and to discover what network traffic is reaching the appliance or leaving the appliance.

The appliance saves the captured packet activity to a file and stores the file locally. You can configure the maximum packet capture file size, how long to run the packet capture, and on which network interface to run the capture. You can also use a filter to limit the packet capture to traffic through a specific port or traffic from a specific client or server IP address.

The **Help and Support > Packet Capture** page on the Security Management appliance displays the list of complete packet capture files stored on the hard drive. When the packet capture process 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.

You can download a packet capture file using the **Download File** button and forward it in an email to Cisco IronPort Customer Support for debugging and troubleshooting purposes. You can also delete a packet capture file by selecting one or more files, and clicking **Delete Selected Files**.

Note In the CLI, use the **packetcapture** command. This command is similar to the **tcpdump** command in UNIX.

Starting a Packet Capture

There are two ways to start a packet capture:

- Starting a Packet Capture from the Command Line Prompt, page 15-3
- Starting a Packet Capture from the GUI, page 15-3

Starting a Packet Capture from the Command Line Prompt

To start a packet capture, type **packetcapture > start** command from the command line prompt. If you need to stop a running packet capture, run the **packetcapture > stop** command. The appliance stops the packet capture when the session ends.

Starting a Packet Capture from the GUI

Procedure

Step 1 On the Security Management appliance, choose **Help and Support > Packet Capture**.

Step 2 Select **Start Capture**.

Step 3 The following figure shows the Packet Capture process running.

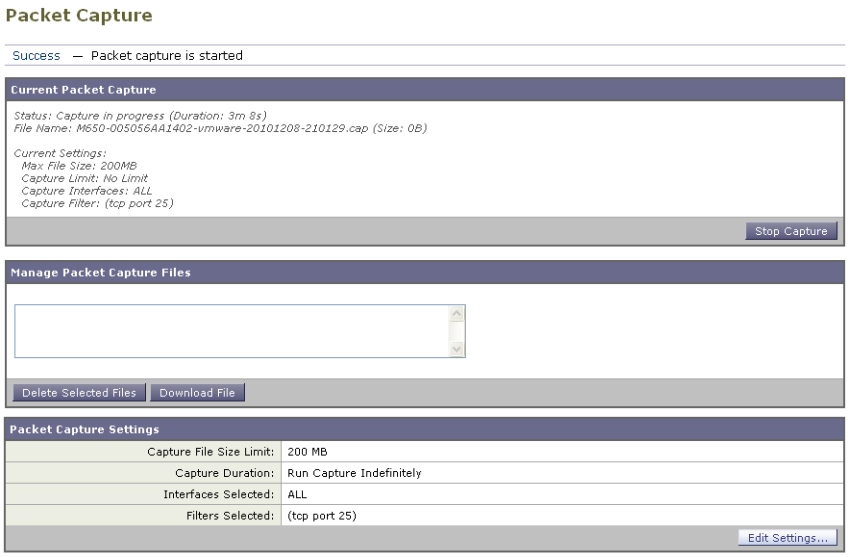

To stop a running capture, click **Stop Capture**. A running capture that has previously been started is preserved between sessions.

Note The GUI only displays packet captures started in the GUI, not from the command line prompt using the **packetcapture > start** command. Similarly, the command line only displays the status of a current packet capture run started from the command line prompt using the **packetcapture > start** command. Only one capture may be running at a time.

Editing Packet Capture Settings

There are two ways to edit packet capture settings:

- Editing a Packet Capture Settings from the Command Line Prompt, page 15-4
- Editing Packet Capture Settings from the GUI, page 15-4

Editing a Packet Capture Settings from the Command Line Prompt

To edit the packet capture settings in the CLI, run the **packetcapture > setup** command from the command line prompt.

Editing Packet Capture Settings from the GUI

Procedure

Step 1 On the Security Management appliance, choose **Help and Support > Packet Capture**.

Step 2 Select **Edit Settings**.

Step 3 Configure settings:

Table 15-1 Packet Capture Configuration Options

Option	Description
Capture file size limit	The maximum file size for all packet capture files in megabytes.
Capture Duration	Choose how long to run the packet capture: • **Run Capture Until File Size Limit Reached**. The packet capture runs until the file size limit is reached. • **Run Capture Until Time Elapsed Reaches**. The packet capture runs until the configured time has passed. You can enter the time in seconds (`s`), minutes (`m`), or hours (`h`). If you enter the amount of time without specifying the units, AsyncOS uses seconds by default. This option is only available in the GUI. **Note** The packet capture file is split into ten parts. If the file reaches the maximum size limit before the entire time has elapsed, 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. • **Run Capture Indefinitely**. The packet capture runs until you manually stop it. **Note** If the file reaches the maximum size limit before you manually stop the packet capture, the oldest part of the file is deleted (the data is discarded) and a new part starts with the current packet capture data. You can always manually stop any packet capture.
Interface	Select the network interface on which to run the packet capture.
Filters	Choose whether or not to apply a filter to the packet capture to reduce the amount of data stored in the packet capture. You can use of the predefined filters to filter by port, client IP, or server IP (GUI only), or you can create a custom filter using any syntax supported by the Unix **tcpdump** command, such as `host 10.10.10.10 && port 80`.

Step 4 Submit your changes. These changes take effect immediately and do not require a commit.

APPENDIX A

IP Interfaces and Accessing the Appliance

You can access any IP interface you create on the appliance through a variety of services.

By default, the following services are either enabled or disabled on each interface:

Table A-1 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 Cisco IronPort 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 *Cisco IronPort AsyncOS for Email Security Advanced Configuration Guide*.

Figure A-1 IP Interfaces Page

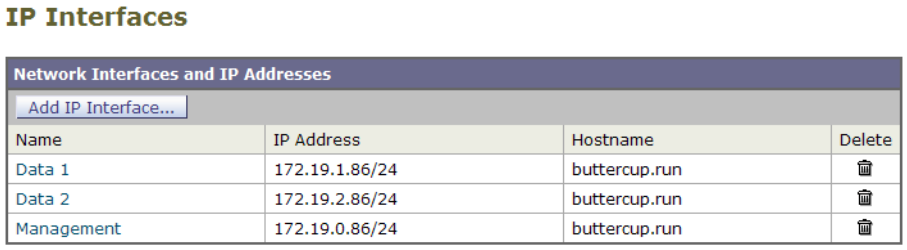

Configuring IP Interfaces

The Management Appliance > Network > IP Interfaces page (and `interfaceconfig` 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 A-2 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, Cisco IronPort 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 Cisco IronPort Spam Quarantine.

Note If you have completed the System Setup Wizard as described in Chapter 2, "Setup, Installation, and Basic Configuration" and committed the changes, the Management interface should already be configured on the appliance.

Creating IP Interfaces Using the GUI

Procedure

Step 1 Choose **Management Appliance > Network > IP Interfaces**.

Step 2 Click **Add IP Interface**.

Step 3 Enter a name for the interface.

Step 4 Select an Ethernet port and enter an IP address.

Step 5 Enter the netmask for the IP address.

Step 6 Enter a hostname for the interface.

Step 7 Select the check box next to each service you want to enable on this IP interface. Change the corresponding port if necessary.

Step 8 Select whether to enable redirecting HTTP to HTTPS for appliance management on the interface.

Step 9 If you are using the Cisco IronPort 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 Cisco IronPort Spam Quarantine, and whether to use the hostname as the URL or provide a custom URL.

Step 10 Submit and commit your changes.

Accessing the Appliance via FTP

Warning **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.**

Procedure

Step 1 Use the Management Appliance > Network > IP Interfaces page (or the `interfaceconfig` command) to enable FTP access for the interface.

In this example, the Management interface is edited to enable FTP access on port 21 (the default port):

Figure A-2 Edit IP Interface Page

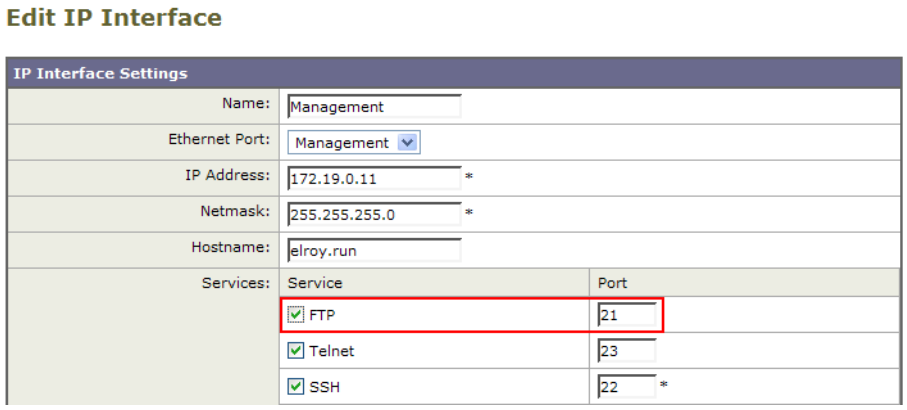

Note Remember to commit your changes before moving on to the next step.

Step 2 Access the interface via FTP. Ensure you are using the correct IP address for the interface.

Example:

```
ftp 192.168.42.42
```

Many browsers also allow you to access interfaces via FTP.

Example:

```
ftp://192.10.10.10
```

Step 3 Browse to the directory for the specific task you are trying to accomplish. After you have accessed an interface via FTP, you can browse the following directories to copy and add ("GET" and "PUT") files. See Table A-3.

Table A-3 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 *Cisco IronPort AsyncOS for Email Security Advanced Configuration Guide* 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

Table A-3 Directories Available for Access (continued)

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 4 Use your FTP program to upload and download files to and from the appropriate directory.

Secure Copy (scp) Access

If your client operating system supports a secure copy (`scp`) command, you can copy files to and from the directories listed in Table A-3 on page A-5. 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 Cisco IronPort 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, page 13-25.

Accessing via a Serial Connection

If you are connecting to the appliance via a serial connection, Figure A-3 illustrates the pin numbers for the serial port connector, and Table A-4 defines the pin assignments and interface signals for the serial port connector.

Figure A-3 Pin Numbers for the Serial Port

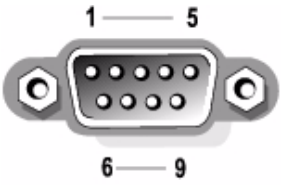

Table A-4 Serial Port Pin Assignments

Pin	Signal	I/O	Definition
1	DCD	I	Data carrier detect
2	SIN	I	Serial input
3	SOUT	O	Serial output
4	DTR	O	Data terminal ready
5	GND	n/a	Signal ground
6	DSR	I	Data set ready
7	RTS	I	Request to send
8	CTS	O	Clear to send
9	RI	I	Ring indicator
Shell	n/a	n/a	Chassis ground

APPENDIX B

Assigning Network and IP Addresses

This appendix describes general rules on networks and IP address assignments, and it presents some strategies for connecting the Cisco IronPort appliance to your network.

Topics included in this appendix include:

- Ethernet Interfaces, page B-1
- Selecting IP Addresses and Netmasks, page B-1
- Strategies for Connecting Your Cisco IronPort Appliance, page B-3

Ethernet Interfaces

The Cisco IronPort X1050, C650, and C350 appliances are equipped with as many as four Ethernet interfaces located on the rear panel of the system, depending on the configuration (whether or not you have the optional optical network interface). They are labeled:

- Management
- Data1
- Data2
- Data3
- Data4

Selecting IP Addresses and Netmasks

When you configure the network, the Cisco IronPort 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 Cisco IronPort appliance, these interface names can represent any two interfaces out of the three Cisco IronPort 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 Cisco IronPort appliance is sent to `192.168.1.11`, there is no way to decide which Ethernet interface should be used to deliver the packet. If the two Ethernet interfaces are connected to two separate physical networks, the packet may be delivered to the incorrect network and never find its destination. The Cisco IronPort 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 Cisco IronPort 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 Cisco IronPort 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 Cisco IronPort appliance must always be able to identify a unique interface over which a packet can be delivered. To make this decision, the Cisco IronPort appliance uses a combination of the packet's destination IP address, and the network and IP address settings of its Ethernet interfaces. The following table summarizes the preceding examples:

	Same Network	Different Network
Same Physical Interface	Allowed	Allowed
Different Physical Interface	Not allowed	Allowed

Strategies for Connecting Your Cisco IronPort Appliance

Keep the following in mind when connecting your Cisco IronPort 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 Cisco IronPort appliance interfaces that you choose to connect and how you address them should be dictated by the complexity of your underlying network. It is not necessary to connect multiple interfaces if your network topology or data volumes do not call for it. It is also possible to keep the connection simple at first as you familiarize yourself with the gateway and then increase the connectivity as volume and network topology require it.

APPENDIX C

Firewall Information

The following table lists the possible ports that may need to be opened for proper operation of the Cisco IronPort appliance (these are the default values).

Table C-1 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.
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.
80	HTTP	In	AsyncOS IPs	HTTP access to the GUI for system monitoring.
80	HTTP	Out	downloads.cisco.com	Service updates, except for AsyncOS upgrades.
80	HTTP	Out	updates.cisco.com	AsyncOS upgrades.
82	HTTP	In	AsyncOS IPs	Used for viewing the Cisco IronPort Spam Quarantine.
83	HTTPS	In	AsyncOS IPs	Used for viewing the Cisco IronPort Spam Quarantine.
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.
110	TCP	Out	POP server	POP authentication for end users for Cisco IronPort Spam Quarantine.
123	UDP	Out	NTP server	NTP if time servers are outside firewall.

Table C-1 Firewall Ports

Default Port	Protocol	In/Out	Hostname	Purpose
143	TCP	Out	IMAP server	IMAP authentication for end users for Cisco IronPort Spam Quarantine.
161	UDP	In	AsyncOS IPs	SNMP queries.
162	UDP	Out	Management station	SNMP traps.
389 3268	LDAP	Out	LDAP servers	LDAP if LDAP directory servers are outside firewall. LDAP authentication for Cisco IronPort Spam Quarantine.
636 3269	LDAPS	Out	LDAPS	LDAPS — ActiveDirectory's global catalog server.
443	TCP	In	AsyncOS IPs	Secure HTTP (`https`) access to the GUI for system monitoring.
443	TCP	Out	update-static.cisco.com	Verify the latest files for the update server.
443	TCP	Out	phonehome.senderbase.org	Receive/send Outbreak Filters.
514	UDP/TCP	Out	Syslog server	Syslog logging.
2222	CCS	In and out	AsyncOS IPs	Cluster Communication Service (for centralized management).
6025	TCP	In	AsyncOS IPs	Send Cisco IronPort Spam Quarantine data to the Security Management appliance if the external Cisco IronPort Spam Quarantine is enabled.

APPENDIX D

Examples

This appendix illustrates and describes a number of common ways to implement the Security Management appliance, and includes the following sections:

- Example 1: Investigating a User, page D-1
- Example 2: Tracking a URL, page D-5
- Example 3: Investigating Top URL Categories Visited, page D-6

Web Security Appliance Examples

This section describes examples using a Security Management appliance and Web Security appliances.

Note All of these scenarios assume that you have enabled web reporting and web tracking on the Security Management appliance *and* on your Web Security appliances. For information on how to enable web tracking and web reporting, see Chapter 5, "Using Centralized Web Reporting and Tracking."

Example 1: Investigating a User

This example demonstrates how a system administrator would investigate a particular user at a company.

In this scenario, a manager has gotten a complaint that an employee is visiting inappropriate web sites at work. To investigate this, the system administrator now needs to track the details of their web activity.

Once the web activity is tracked, a web report is generated with information about the employee's browsing history.

Step 1 On the Security Management appliance, choose **Web > Reporting > Users**.

Step 2 In the **Users** table, click on the **User ID** or **Client IP address** you want to investigate.

If you do not know the User ID or the Client IP address, type what you can remember of the User ID or Client IP address in text field, and click on **Find User ID or Client IP address**. The IP address does not need to be an exact match to return results. The Users table is populated with the User ID and Client IP addresses that you have specified. In this example, we are looking for information on Client IP address 10.251.60.24.

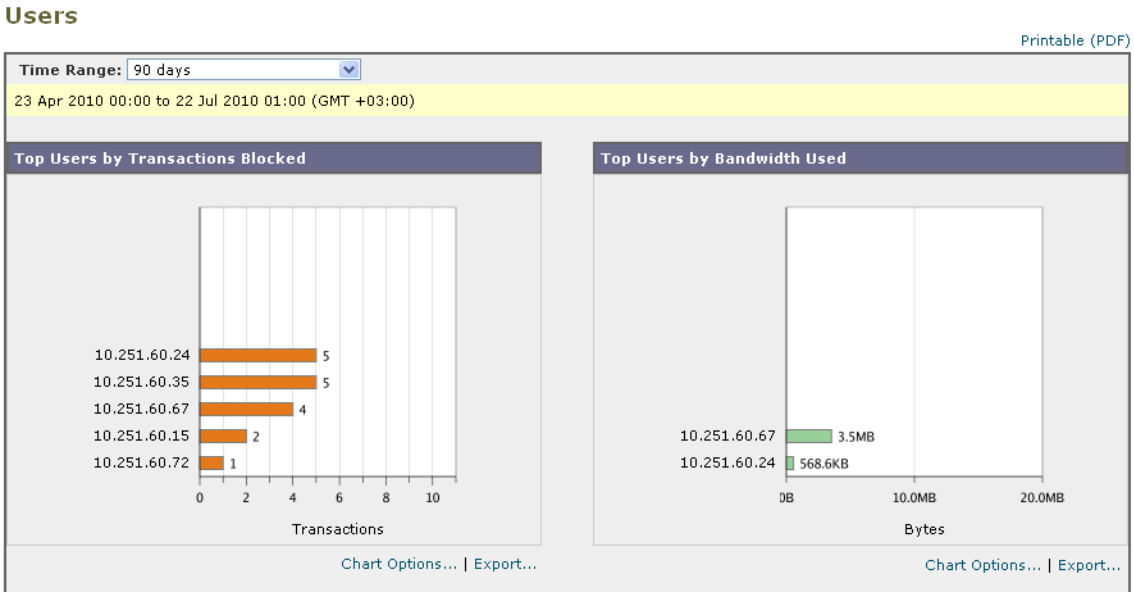

Step 3 Click on IP address **10.251.60.24**.

The User Details page appears for 10.251.60.24.

Users > 10.251.60.24

Printable (PDF)

Time Range: 90 days

31 Aug 2011 00:00 to 29 Nov 2011 15:00 (GMT -08:00)

URL Categories by Total Transactions

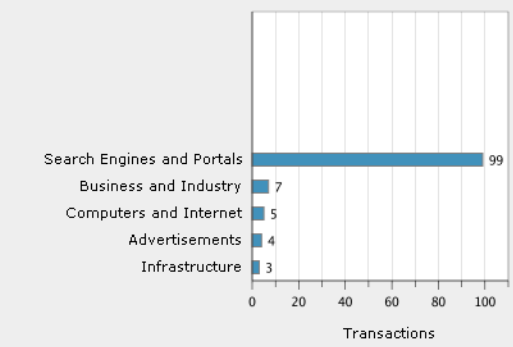

Chart Options... | Export...

Trend by Total Transactions

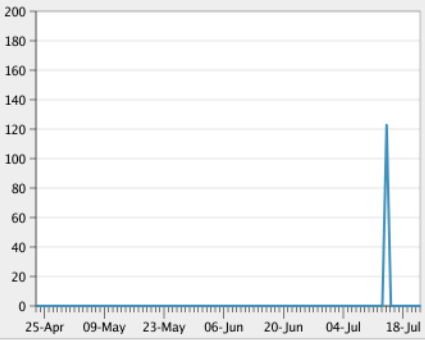

Chart Options... | Export...

URL Categories Matched

URL Category	Bandwidth Used	Time Spent	Blocked URL Category	Transactions Completed	Total Transactions
Search Engines and Portals	447.4KB	00:21	0	99	99
Business and Industry	15.5KB	00:06	0	7	7
Computers and Internet	84.4KB	00:06	0	5	5
Advertisements	16.9KB	00:00	0	4	4
Infrastructure	4,540B	00:00	0	3	3
Totals (all available data):	568.6KB	00:33	0	118	118

Find URL Category

Columns... | Export...

Domains Matched

Domain or IP	Bandwidth Used	Time Spent	Transactions Completed	Transactions Blocked	Total Transactions
[redacted].com	464.5KB	00:36	101	5	106
[redacted].com	83.1KB	00:06	4	0	4
[redacted].com	8,272B	00:00	4	0	4
[redacted].net	15.1KB	00:00	3	0	3
[redacted].com	6,391B	00:06	2	0	2
[redacted].com	1,365B	00:00	1	0	1
[redacted].com	1,231B	00:00	1	0	1
[redacted].com	1,847B	00:00	1	0	1
[redacted]	2,021B	00:00	1	0	1

Find Domain or IP

Columns... | Export...

Applications Matched

No data was found in the selected time range

Malware Threats Detected

No data was found in the selected time range

Policies Matched

Policy Name	Policy Type	Bandwidth Used	Completed Transactions	Blocked Transactions	Total Transactions

From the User Details page you can determine the URL Categories by Total Transactions, Trend by Total Transaction, URL Categories Matched, Domains Matched, Applications Matched, Malware Threats Detected, and Policies Matched.

These categories allow you to find out if, for example, user 10.251.60.24 was trying to access blocked URLs, which could be viewed in the Transactions Blocked column under the Domains section on the page.

Step 4 Click **Export** under the Domains Matched table to view the entire list of Domains and URLs that the user tried to access.

Figure D-1 shows you an exported list of information that has been exported from the user.

Figure D-1 Sample Export Data

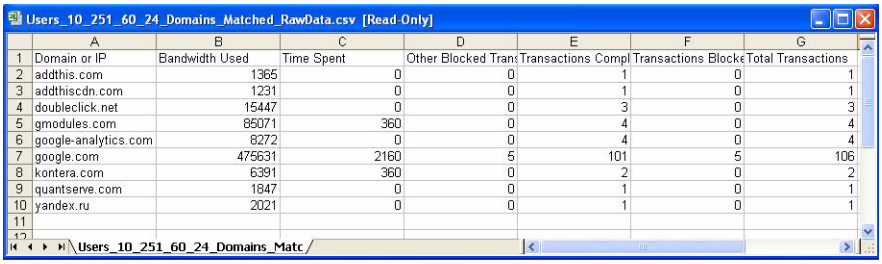
Users_10_251_60_24_Domains_Matched_RawData.csv [Read-Only]

	A	B	C	D	E	F	G
1	Domain or IP	Bandwidth Used	Time Spent	Other Blocked Trans	Transactions Compl	Transactions Blocke	Total Transactions
2	addthis.com	1365	0	0	1	0	1
3	addthiscdn.com	1231	0	0	1	0	1
4	doubleclick.net	15447	0	0	3	0	3
5	gmodules.com	85071	360	0	4	0	4
6	google-analytics.com	8272	0	0	4	0	4
7	google.com	475631	2160	5	101	5	106
8	kontera.com	6391	360	0	2	0	2
9	quantserve.com	1847	0	0	1	0	1
10	yandex.ru	2021	0	0	1	0	1
11							

Users_10_251_60_24_Domains_Matc

From here you can use the Web Tracking feature to track and view this specific user's web usage.

Note It is important to remember that web reporting allows you to retrieve all the domain information that a user goes to, not necessarily the specific URL that is accessed. For information on a specific URL that the user is accessing, what time they went to that URL, whether that URL is allowed, etc., use the Proxy Services tab on the Web Tracking page.

Step 5 Choose **Web > Reporting > Web Tracking**.

Step 6 Click the **Proxy Services** tab.

Step 7 In the User/Client IP Address text field type in the user name or IP address.

In this example we are searching for web tracking information for user 10.251.60.24.

The search results appear.

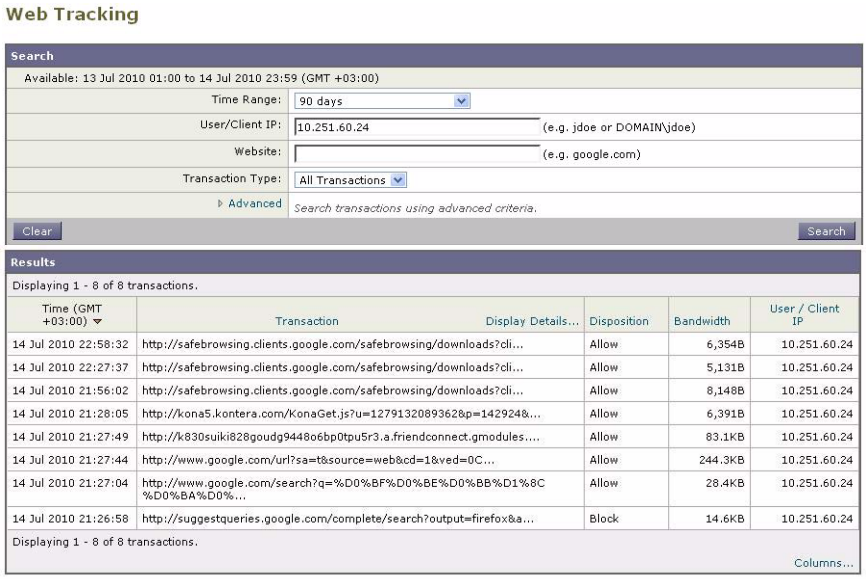

From this page you can view a full list of transactions and URLs visited by the user of the computer that is assigned to the IP Address 10.251.60.24.

Related Topics

Table D-1 lists each of the topics discussed in this example. Click on the link for details on each topic.

Table D-1 Related Topics for Investigating a User

Feature Name	Feature Information
User Page	Users Report (Web), page 5-16
User Details Page	User Details (Web Reporting), page 5-19
Exporting Report Data	Printing and Exporting Reporting and Tracking Data, page 3-8
Proxy Services tab on the Web Tracking page	Searching for Transactions Processed by Web Proxy Services, page 5-51

Example 2: Tracking a URL

In this scenario, a Sales manager wants to find out what the top five visited web sites are at their company are for the last week. Additionally, the manager wants to know which users are going to those websites.

Step 1 On the Security Management appliance, choose **Web > Reporting > Web Sites**.

Step 2 From the Time Range drop-down list, choose **Week**.

Step 3 Scroll down to the Domains section to view the domains, or web sites that have been visited.

The top 25 web sites that have been visited will be displayed in the Domains Matched table. In the same table you can click on the link in the Domain or IP column to view the actual web sites for a particular address or user.

Related Topics

Table D-2 lists each of the topics discussed in this example. Click on the link for details on each topic.

Table D-2 Related Topics for Tracking a URL

Feature Name	Feature Information
Web Sites Page	Web Sites Report, page 5-22

Example 3: Investigating Top URL Categories Visited

In this scenario, the Human Resources manager wants to know what the top three URL categories her employees are visiting over the 30 days. Additionally, a network manager wants to get this information to monitor bandwidth usage, to find out what URLs are taking up the most bandwidth on her network.

The example below is to show how you can gather data for several people covering several points of interest, while only having to generate one report.

Step 1 On the Security Management appliance, choose **Web > Reporting > URL Categories**.

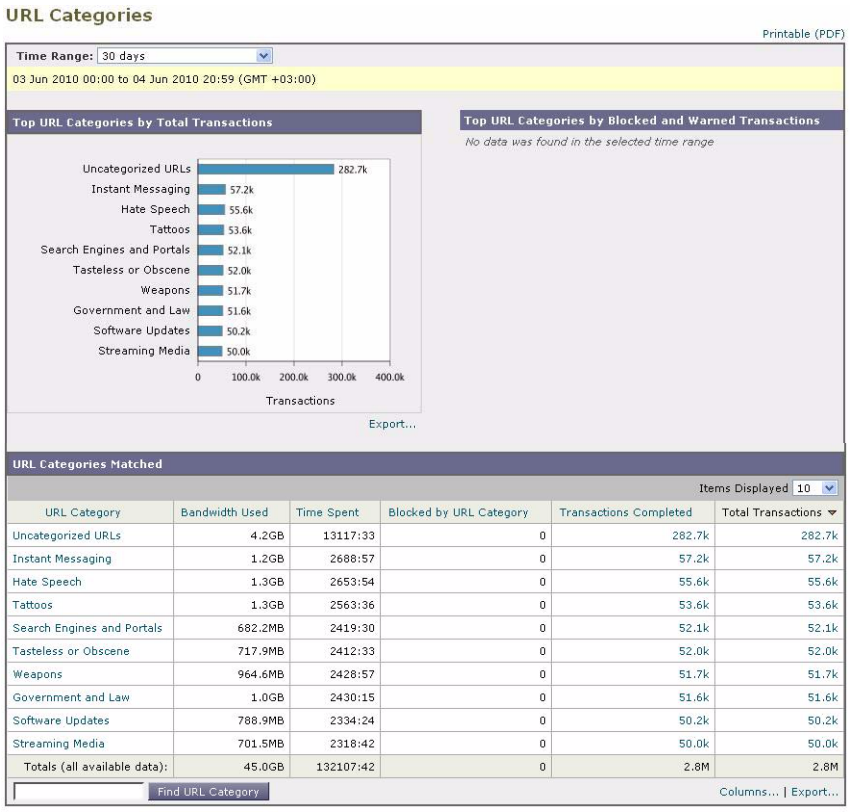

From the URL Categories page in this example, you can see that of the top 10 URL Categories by Total Transactions graph reveals, there were 282 K of Uncategorized URLs that were accessed, as well as Instant Messaging, Hate Speech and Tattoo sites, and so forth.

At this point you can export that raw data to an Excel spreadsheet, by clicking the **Export** link and send this file to the Human Resources manager. But remember, your network manager wants to know the bandwidth usage by each URL.

Step 2 Scroll down to the **URL Categories Matched** table, to view the Bandwidth Used column.

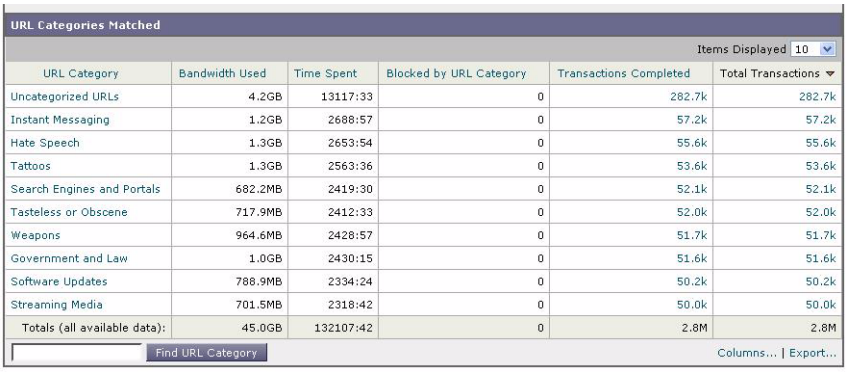

From the **URL Categories Matched** table, you can see the Bandwidth Usage for all of the URL Categories. Again, you can click the **Export** link and send this file to the Network manager. For finer granularity though, click on the Instant Messaging link to find out which users are taking up the bandwidth. The following page appears.

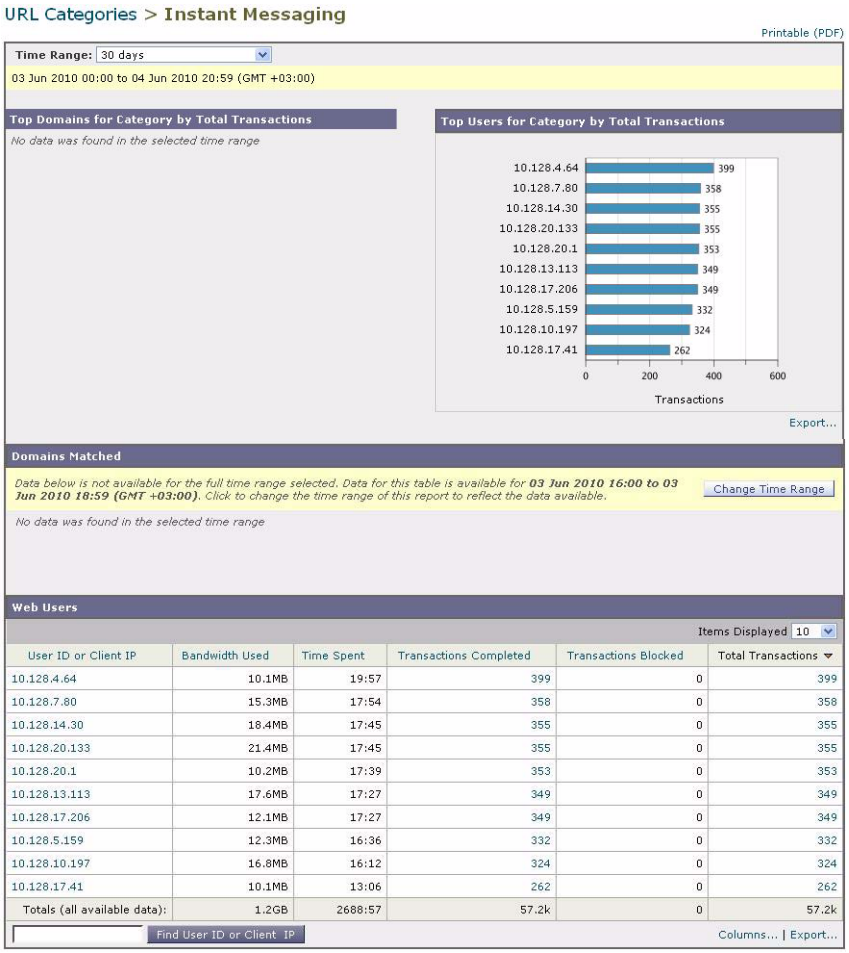

From this page, the network manager can see the top 10 users for Instant Messaging sites.

This pages reveals that in the last 30 days, user 10.128.4.64 has spent 19 hours and 57 minutes on an Instant Messaging site; and the bandwidth usage for this time was 10.1 MB.

Related Topics

Table D-3 lists each of the topics discussed in this example. Click on the link for details on each topic.

Table D-3 Related Topics for Investigating the Top URL Categories

Feature Name	Feature Information
URL Categories Page	URL Categories Report, page 5-24
Exporting Report Data	Printing and Exporting Reporting and Tracking Data, page 3-8

APPENDIX E

End User License Agreement

Cisco Systems End User License Agreement

IMPORTANT: PLEASE READ THIS END USER LICENSE AGREEMENT CAREFULLY. IT IS VERY IMPORTANT THAT YOU CHECK THAT YOU ARE PURCHASING CISCO SOFTWARE OR EQUIPMENT FROM AN APPROVED SOURCE AND THAT YOU, OR THE ENTITY YOU REPRESENT (COLLECTIVELY, THE "CUSTOMER") HAVE BEEN REGISTERED AS THE END USER FOR THE PURPOSES OF THIS CISCO END USER LICENSE AGREEMENT. IF YOU ARE NOT REGISTERED AS THE END USER YOU HAVE NO LICENSE TO USE THE SOFTWARE AND THE LIMITED WARRANTY IN THIS END USER LICENSE AGREEMENT DOES NOT APPLY. ASSUMING YOU HAVE PURCHASED FROM AN APPROVED SOURCE, DOWNLOADING, INSTALLING OR USING CISCO OR CISCO-SUPPLIED SOFTWARE CONSTITUTES ACCEPTANCE OF THIS AGREEMENT.

CISCO SYSTEMS, INC. OR ITS SUBSIDIARY LICENSING THE SOFTWARE INSTEAD OF CISCO SYSTEMS, INC. ("CISCO") IS WILLING TO LICENSE THIS SOFTWARE TO YOU ONLY UPON THE CONDITION THAT YOU PURCHASED THE SOFTWARE FROM AN APPROVED SOURCE AND THAT YOU ACCEPT ALL OF THE TERMS CONTAINED IN THIS END USER LICENSE AGREEMENT PLUS ANY ADDITIONAL LIMITATIONS ON THE LICENSE SET FORTH IN A SUPPLEMENTAL LICENSE AGREEMENT ACCOMPANYING THE PRODUCT OR AVAILABLE AT THE TIME OF YOUR ORDER (COLLECTIVELY THE "AGREEMENT"). TO THE EXTENT OF ANY CONFLICT BETWEEN THE TERMS OF THIS END USER LICENSE AGREEMENT AND ANY SUPPLEMENTAL LICENSE AGREEMENT, THE SUPPLEMENTAL LICENSE AGREEMENT SHALL APPLY. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE, YOU ARE REPRESENTING THAT YOU PURCHASED THE SOFTWARE FROM AN APPROVED SOURCE AND BINDING YOURSELF TO THE AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THE AGREEMENT, THEN CISCO IS UNWILLING TO LICENSE THE SOFTWARE TO YOU AND (A) YOU MAY NOT DOWNLOAD, INSTALL OR USE THE SOFTWARE, AND (B) YOU MAY RETURN THE SOFTWARE (INCLUDING ANY UNOPENED CD PACKAGE AND ANY WRITTEN MATERIALS) FOR A FULL REFUND, OR, IF THE SOFTWARE AND WRITTEN MATERIALS ARE SUPPLIED AS PART OF ANOTHER PRODUCT, YOU MAY RETURN THE ENTIRE PRODUCT FOR A FULL REFUND. YOUR RIGHT TO RETURN AND REFUND EXPIRES 30 DAYS AFTER PURCHASE FROM AN APPROVED SOURCE, AND APPLIES ONLY IF YOU ARE THE ORIGINAL AND REGISTERED END USER PURCHASER. FOR THE PURPOSES OF THIS END USER LICENSE AGREEMENT, AN "APPROVED SOURCE" MEANS (A) CISCO; OR (B) A DISTRIBUTOR OR SYSTEMS INTEGRATOR AUTHORIZED BY CISCO TO DISTRIBUTE / SELL CISCO EQUIPMENT, SOFTWARE AND SERVICES WITHIN YOUR TERRITORY TO END USERS; OR (C) A RESELLER AUTHORIZED BY ANY SUCH DISTRIBUTOR OR SYSTEMS

INTEGRATOR IN ACCORDANCE WITH THE TERMS OF THE DISTRIBUTOR'S AGREEMENT WITH CISCO TO DISTRIBUTE / SELL THE CISCO EQUIPMENT, SOFTWARE AND SERVICES WITHIN YOUR TERRITORY TO END USERS.

THE FOLLOWING TERMS OF THE AGREEMENT GOVERN CUSTOMER'S USE OF THE SOFTWARE (DEFINED BELOW), EXCEPT TO THE EXTENT: (A) THERE IS A SEPARATE SIGNED CONTRACT BETWEEN CUSTOMER AND CISCO GOVERNING CUSTOMER'S USE OF THE SOFTWARE, OR (B) THE SOFTWARE INCLUDES A SEPARATE "CLICK-ACCEPT" LICENSE AGREEMENT OR THIRD PARTY LICENSE AGREEMENT AS PART OF THE INSTALLATION OR DOWNLOAD PROCESS GOVERNING CUSTOMER'S USE OF THE SOFTWARE. TO THE EXTENT OF A CONFLICT BETWEEN THE PROVISIONS OF THE FOREGOING DOCUMENTS, THE ORDER OF PRECEDENCE SHALL BE (1)THE SIGNED CONTRACT, (2) THE CLICK-ACCEPT AGREEMENT OR THIRD PARTY LICENSE AGREEMENT, AND (3) THE AGREEMENT. FOR PURPOSES OF THE AGREEMENT, "SOFTWARE" SHALL MEAN COMPUTER PROGRAMS, INCLUDING FIRMWARE AND COMPUTER PROGRAMS EMBEDDED IN CISCO EQUIPMENT, AS PROVIDED TO CUSTOMER BY AN APPROVED SOURCE, AND ANY UPGRADES, UPDATES, BUG FIXES OR MODIFIED VERSIONS THERETO (COLLECTIVELY, "UPGRADES"), ANY OF THE SAME WHICH HAS BEEN RELICENSED UNDER THE CISCO SOFTWARE TRANSFER AND RE-LICENSING POLICY (AS MAY BE AMENDED BY CISCO FROM TIME TO TIME) OR BACKUP COPIES OF ANY OF THE FOREGOING.

License. Conditioned upon compliance with the terms and conditions of the Agreement, Cisco grants to Customer a nonexclusive and nontransferable license to use for Customer's internal business purposes the Software and the Documentation for which Customer has paid the required license fees to an Approved Source. "Documentation" means written information (whether contained in user or technical manuals, training materials, specifications or otherwise) pertaining to the Software and made available by an Approved Source with the Software in any manner (including on CD-Rom, or on-line). In order to use the Software, Customer may be required to input a registration number or product authorization key and register Customer's copy of the Software online at Cisco's website to obtain the necessary license key or license file.

Customer's license to use the Software shall be limited to, and Customer shall not use the Software in excess of, a single hardware chassis or card or such other limitations as are set forth in the applicable Supplemental License Agreement or in the applicable purchase order which has been accepted by an Approved Source and for which Customer has paid to an Approved Source the required license fee (the "Purchase Order").

Unless otherwise expressly provided in the Documentation or any applicable Supplemental License Agreement, Customer shall use the Software solely as embedded in, for execution on, or (where the applicable Documentation permits installation on non-Cisco equipment) for communication with Cisco equipment owned or leased by Customer and used for Customer's internal business purposes. No other licenses are granted by implication, estoppel or otherwise.

For evaluation or beta copies for which Cisco does not charge a license fee, the above requirement to pay license fees does not apply.

General Limitations. This is a license, not a transfer of title, to the Software and Documentation, and Cisco retains ownership of all copies of the Software and Documentation. Customer acknowledges that the Software and Documentation contain trade secrets of Cisco or its suppliers or licensors, including but not limited to the specific internal design and structure of individual programs and associated interface information. Except as otherwise expressly provided under the Agreement, Customer shall only use the Software in connection with the use of Cisco equipment purchased by the Customer from an Approved Source and Customer shall have no right, and Customer specifically agrees not to:

(i) transfer, assign or sublicense its license rights to any other person or entity (other than in compliance with any Cisco relicensing/transfer policy then in force), or use the Software on Cisco equipment not purchased by the Customer from an Approved Source or on secondhand Cisco equipment, and Customer acknowledges that any attempted transfer, assignment, sublicense or use shall be void;

(ii) make error corrections to or otherwise modify or adapt the Software or create derivative works based upon the Software, or permit third parties to do the same;

(iii) reverse engineer or decompile, decrypt, disassemble or otherwise reduce the Software to human-readable form, except to the extent otherwise expressly permitted under applicable law notwithstanding this restriction or except to the extent that Cisco is legally required to permit such specific activity pursuant to any applicable open source license;

(iv) publish any results of benchmark tests run on the Software;

(v) use or permit the Software to be used to perform services for third parties, whether on a service bureau or time sharing basis or otherwise, without the express written authorization of Cisco; or

(vi) disclose, provide, or otherwise make available trade secrets contained within the Software and Documentation in any form to any third party without the prior written consent of Cisco. Customer shall implement reasonable security measures to protect such trade secrets.

To the extent required by applicable law, and at Customer's written request, Cisco shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of Cisco's applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Cisco makes such information available.

Software, Upgrades and Additional Copies. NOTWITHSTANDING ANY OTHER PROVISION OF THE AGREEMENT: (1) CUSTOMER HAS NO LICENSE OR RIGHT TO MAKE OR USE ANY ADDITIONAL COPIES OR UPGRADES UNLESS CUSTOMER, AT THE TIME OF MAKING OR ACQUIRING SUCH COPY OR UPGRADE, ALREADY HOLDS A VALID LICENSE TO THE ORIGINAL SOFTWARE AND HAS PAID THE APPLICABLE FEE TO AN APPROVED SOURCE FOR THE UPGRADE OR ADDITIONAL COPIES; (2) USE OF UPGRADES IS LIMITED TO CISCO EQUIPMENT SUPPLIED BY AN APPROVED SOURCE FOR WHICH CUSTOMER IS THE ORIGINAL END USER PURCHASER OR LESSEE OR OTHERWISE HOLDS A VALID LICENSE TO USE THE SOFTWARE WHICH IS BEING UPGRADED; AND (3) THE MAKING AND USE OF ADDITIONAL COPIES IS LIMITED TO NECESSARY BACKUP PURPOSES ONLY.

Proprietary Notices. Customer agrees to maintain and reproduce all copyright, proprietary, and other notices on all copies, in any form, of the Software in the same form and manner that such copyright and other proprietary notices are included on the Software. Except as expressly authorized in the Agreement, Customer shall not make any copies or duplicates of any Software without the prior written permission of Cisco.

Term and Termination. The Agreement and the license granted herein shall remain effective until terminated. Customer may terminate the Agreement and the license at any time by destroying all copies of Software and any Documentation. Customer's rights under the Agreement will terminate immediately without notice from Cisco if Customer fails to comply with any provision of the Agreement. Upon termination, Customer shall destroy all copies of Software and Documentation in its possession or control. All confidentiality obligations of Customer, all restrictions and limitations imposed on the Customer under the section titled "General Limitations" and all limitations of liability and disclaimers and restrictions of warranty shall survive termination of this Agreement. In addition, the provisions of the sections titled "U.S. Government End User Purchasers" and "General Terms Applicable to the Limited Warranty Statement and End User License Agreement" shall survive termination of the Agreement.

Customer Records. Customer grants to Cisco and its independent accountants the right to examine Customer's books, records and accounts during Customer's normal business hours to verify compliance with this Agreement. In the event such audit discloses non-compliance with this Agreement, Customer shall promptly pay to Cisco the appropriate license fees, plus the reasonable cost of conducting the audit.

Export, Re-Export, Transfer and Use Controls. The Software, Documentation and technology or direct products thereof (hereafter referred to as Software and Technology), supplied by Cisco under the Agreement are subject to export controls under the laws and regulations of the United States (U.S.) and any other applicable countries' laws and regulations. Customer shall comply with such laws and regulations governing export, re-export, transfer and use of Cisco Software and Technology and will obtain all required U.S. and local authorizations, permits, or licenses. Cisco and Customer each agree to provide the other information, support documents, and assistance as may reasonably be required by the other in connection with securing authorizations or licenses. Information regarding compliance with export, re-export, transfer and use may be located at the following URL:

http://www.cisco.com/web/about/doing_business/legal/global_export_trade/general_export/contract_compliance.html.

U.S. Government End User Purchasers. The Software and Documentation qualify as "commercial items," as that term is defined at Federal Acquisition Regulation ("FAR") (48 C.F.R.) 2.101, consisting of "commercial computer software" and "commercial computer software documentation" as such terms are used in FAR 12.212. Consistent with FAR 12.212 and DoD FAR Supp. 227.7202-1 through 227.7202-4, and notwithstanding any other FAR or other contractual clause to the contrary in any agreement into which the Agreement may be incorporated, Customer may provide to Government end user or, if the Agreement is direct, Government end user will acquire, the Software and Documentation with only those rights set forth in the Agreement. Use of either the Software or Documentation or both constitutes agreement by the Government that the Software and Documentation are "commercial computer software" and "commercial computer software documentation," and constitutes acceptance of the rights and restrictions herein.

Identified Components; Additional Terms. The Software may contain or be delivered with one or more components, which may include third-party components, identified by Cisco in the Documentation, readme.txt file, third-party click-accept or elsewhere (e.g. on www.cisco.com) (the "Identified Component(s)") as being subject to different license agreement terms, disclaimers of warranties, limited warranties or other terms and conditions (collectively, "Additional Terms") than those set forth herein. You agree to the applicable Additional Terms for any such Identified Component(s)."

Limited Warranty

Subject to the limitations and conditions set forth herein, Cisco warrants that commencing from the date of shipment to Customer (but in case of resale by an Approved Source other than Cisco, commencing not more than ninety (90) days after original shipment by Cisco), and continuing for a period of the longer of (a) ninety (90) days or (b) the warranty period (if any) expressly set forth as applicable specifically to software in the warranty card accompanying the product of which the Software is a part (the "Product") (if any): (a) the media on which the Software is furnished will be free of defects in materials and workmanship under normal use; and (b) the Software substantially conforms to the Documentation. The date of shipment of a Product by Cisco is set forth on the packaging material in which the Product is shipped. Except for the foregoing, the Software is provided "AS IS". This limited warranty extends only to the Software purchased from an Approved Source by a Customer who is the first registered end user. Customer's sole and exclusive remedy and the entire liability of Cisco and its suppliers under this limited warranty will be (i) replacement of defective media and/or (ii) at Cisco's option, repair, replacement, or refund of the purchase price of the Software, in both cases subject to the condition that any error or defect constituting a breach of this limited warranty is reported to the Approved Source supplying the Software to Customer, within the warranty period. Cisco or the Approved Source supplying the Software to Customer may, at its option, require return of the Software and/or Documentation as a condition to the remedy. In no event does Cisco warrant that the Software is

error free or that Customer will be able to operate the Software without problems or interruptions. In addition, due to the continual development of new techniques for intruding upon and attacking networks, Cisco does not warrant that the Software or any equipment, system or network on which the Software is used will be free of vulnerability to intrusion or attack.

Restrictions. This warranty does not apply if the Software, Product or any other equipment upon which the Software is authorized to be used (a) has been altered, except by Cisco or its authorized representative, (b) has not been installed, operated, repaired, or maintained in accordance with instructions supplied by Cisco, (c) has been subjected to abnormal physical or electrical stress, abnormal environmental conditions, misuse, negligence, or accident; or (d) is licensed for beta, evaluation, testing or demonstration purposes. The Software warranty also does not apply to (e) any temporary Software modules; (f) any Software not posted on Cisco's Software Center; (g) any Software that Cisco expressly provides on an "AS IS" basis on Cisco's Software Center; (h) any Software for which an Approved Source does not receive a license fee; and (i) Software supplied by any third party which is not an Approved Source.

DISCLAIMER OF WARRANTY

EXCEPT AS SPECIFIED IN THIS WARRANTY SECTION, ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS, AND WARRANTIES INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTY OR CONDITION OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, SATISFACTORY QUALITY, NON-INTERFERENCE, ACCURACY OF INFORMATIONAL CONTENT, OR ARISING FROM A COURSE OF DEALING, LAW, USAGE, OR TRADE PRACTICE, ARE HEREBY EXCLUDED TO THE EXTENT ALLOWED BY APPLICABLE LAW AND ARE EXPRESSLY DISCLAIMED BY CISCO, ITS SUPPLIERS AND LICENSORS. TO THE EXTENT THAT ANY OF THE SAME CANNOT BE EXCLUDED, SUCH IMPLIED CONDITION, REPRESENTATION AND/OR WARRANTY IS LIMITED IN DURATION TO THE EXPRESS WARRANTY PERIOD REFERRED TO IN THE "LIMITED WARRANTY" SECTION ABOVE. BECAUSE SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATIONS ON HOW LONG AN IMPLIED WARRANTY LASTS, THE ABOVE LIMITATION MAY NOT APPLY IN SUCH STATES. THIS WARRANTY GIVES CUSTOMER SPECIFIC LEGAL RIGHTS, AND CUSTOMER MAY ALSO HAVE OTHER RIGHTS WHICH VARY FROM JURISDICTION TO JURISDICTION. This disclaimer and exclusion shall apply even if the express warranty set forth above fails of its essential purpose.

Disclaimer of Liabilities - Limitation of Liability. IF YOU ACQUIRED THE SOFTWARE IN THE UNITED STATES, LATIN AMERICA, CANADA, JAPAN OR THE CARIBBEAN, NOTWITHSTANDING ANYTHING ELSE IN THE AGREEMENT TO THE CONTRARY, ALL LIABILITY OF CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS COLLECTIVELY, TO CUSTOMER, WHETHER IN CONTRACT, TORT (INCLUDING NEGLIGENCE), BREACH OF WARRANTY OR OTHERWISE, SHALL NOT EXCEED THE PRICE PAID BY CUSTOMER TO ANY APPROVED SOURCE FOR THE SOFTWARE THAT GAVE RISE TO THE CLAIM OR IF THE SOFTWARE IS PART OF ANOTHER PRODUCT, THE PRICE PAID FOR SUCH OTHER PRODUCT. THIS LIMITATION OF LIABILITY FOR SOFTWARE IS CUMULATIVE AND NOT PER INCIDENT (I.E. THE EXISTENCE OF TWO OR MORE CLAIMS WILL NOT ENLARGE THIS LIMIT).

IF YOU ACQUIRED THE SOFTWARE IN EUROPE, THE MIDDLE EAST, AFRICA, ASIA OR OCEANIA, NOTWITHSTANDING ANYTHING ELSE IN THE AGREEMENT TO THE CONTRARY, ALL LIABILITY OF CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS COLLECTIVELY, TO CUSTOMER, WHETHER IN CONTRACT, TORT (INCLUDING NEGLIGENCE), BREACH OF WARRANTY OR OTHERWISE, SHALL NOT EXCEED THE PRICE PAID BY CUSTOMER TO CISCO FOR THE SOFTWARE THAT GAVE RISE TO THE CLAIM OR IF THE SOFTWARE IS PART OF ANOTHER PRODUCT, THE PRICE PAID

FOR SUCH OTHER PRODUCT. THIS LIMITATION OF LIABILITY FOR SOFTWARE IS CUMULATIVE AND NOT PER INCIDENT (I.E. THE EXISTENCE OF TWO OR MORE CLAIMS WILL NOT ENLARGE THIS LIMIT). NOTHING IN THE AGREEMENT SHALL LIMIT (I) THE LIABILITY OF CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS TO CUSTOMER FOR PERSONAL INJURY OR DEATH CAUSED BY THEIR NEGLIGENCE, (II) CISCO'S LIABILITY FOR FRAUDULENT MISREPRESENTATION, OR (III) ANY LIABILITY OF CISCO WHICH CANNOT BE EXCLUDED UNDER APPLICABLE LAW.

Disclaimer of Liabilities - Waiver of Consequential Damages and Other Losses. IF YOU ACQUIRED THE SOFTWARE IN THE UNITED STATES, LATIN AMERICA, THE CARIBBEAN OR CANADA, REGARDLESS OF WHETHER ANY REMEDY SET FORTH HEREIN FAILS OF ITS ESSENTIAL PURPOSE OR OTHERWISE, IN NO EVENT WILL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY LOST REVENUE, PROFIT, OR LOST OR DAMAGED DATA, BUSINESS INTERRUPTION, LOSS OF CAPITAL, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY OR WHETHER ARISING OUT OF THE USE OF OR INABILITY TO USE SOFTWARE OR OTHERWISE AND EVEN IF CISCO OR ITS SUPPLIERS OR LICENSORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATION OR EXCLUSION OF CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATION MAY NOT APPLY TO YOU.

IF YOU ACQUIRED THE SOFTWARE IN JAPAN, EXCEPT FOR LIABILITY ARISING OUT OF OR IN CONNECTION WITH DEATH OR PERSONAL INJURY, FRAUDULENT MISREPRESENTATION, AND REGARDLESS OF WHETHER ANY REMEDY SET FORTH HEREIN FAILS OF ITS ESSENTIAL PURPOSE OR OTHERWISE, IN NO EVENT WILL CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS BE LIABLE FOR ANY LOST REVENUE, PROFIT, OR LOST OR DAMAGED DATA, BUSINESS INTERRUPTION, LOSS OF CAPITAL, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY OR WHETHER ARISING OUT OF THE USE OF OR INABILITY TO USE SOFTWARE OR OTHERWISE AND EVEN IF CISCO OR ANY APPROVED SOURCE OR THEIR SUPPLIERS OR LICENSORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

IF YOU ACQUIRED THE SOFTWARE IN EUROPE, THE MIDDLE EAST, AFRICA, ASIA OR OCEANIA, IN NO EVENT WILL CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS, BE LIABLE FOR ANY LOST REVENUE, LOST PROFIT, OR LOST OR DAMAGED DATA, BUSINESS INTERRUPTION, LOSS OF CAPITAL, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES, HOWSOEVER ARISING, INCLUDING, WITHOUT LIMITATION, IN CONTRACT, TORT (INCLUDING NEGLIGENCE) OR WHETHER ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF, IN EACH CASE, CISCO, ITS AFFILIATES, OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS AND LICENSORS, HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATION OR EXCLUSION OF CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATION MAY NOT FULLY APPLY TO YOU. THE FOREGOING EXCLUSION SHALL NOT APPLY TO ANY LIABILITY ARISING OUT OF OR IN CONNECTION WITH: (I) DEATH OR PERSONAL INJURY, (II) FRAUDULENT MISREPRESENTATION, OR (III) CISCO'S LIABILITY IN CONNECTION WITH ANY TERMS THAT CANNOT BE EXCLUDED UNDER APPLICABLE LAW.

Customer acknowledges and agrees that Cisco has set its prices and entered into the Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the parties.

Controlling Law, Jurisdiction. If you acquired, by reference to the address on the purchase order accepted by the Approved Source, the Software in the United States, Latin America, or the Caribbean, the Agreement and warranties ("Warranties") are controlled by and construed under the laws of the State of California, United States of America, notwithstanding any conflicts of law provisions; and the state and federal courts of California shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in Canada, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of the Province of Ontario, Canada, notwithstanding any conflicts of law provisions; and the courts of the Province of Ontario shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in Europe, the Middle East, Africa, Asia or Oceania (excluding Australia), unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of England, notwithstanding any conflicts of law provisions; and the English courts shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. In addition, if the Agreement is controlled by the laws of England, no person who is not a party to the Agreement shall be entitled to enforce or take the benefit of any of its terms under the Contracts (Rights of Third Parties) Act 1999. If you acquired the Software in Japan, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of Japan, notwithstanding any conflicts of law provisions; and the Tokyo District Court of Japan shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in Australia, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of the State of New South Wales, Australia, notwithstanding any conflicts of law provisions; and the State and federal courts of New South Wales shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties. If you acquired the Software in any other country, unless expressly prohibited by local law, the Agreement and Warranties are controlled by and construed under the laws of the State of California, United States of America, notwithstanding any conflicts of law provisions; and the state and federal courts of California shall have exclusive jurisdiction over any claim arising under the Agreement or Warranties.

For all countries referred to above, the parties specifically disclaim the application of the UN Convention on Contracts for the International Sale of Goods. Notwithstanding the foregoing, either party may seek interim injunctive relief in any court of appropriate jurisdiction with respect to any alleged breach of such party's intellectual property or proprietary rights. If any portion hereof is found to be void or unenforceable, the remaining provisions of the Agreement and Warranties shall remain in full force and effect. Except as expressly provided herein, the Agreement constitutes the entire agreement between the parties with respect to the license of the Software and Documentation and supersedes any conflicting or additional terms contained in any Purchase Order or elsewhere, all of which terms are excluded. The Agreement has been written in the English language, and the parties agree that the English version will govern.

Product warranty terms and other information applicable to Cisco products are available at the following URL:

http://www.cisco.com/go/warranty

Supplemental End User License Agreement for Cisco Systems Content Security Software

IMPORTANT: READ CAREFULLY

This Supplemental End User License Agreement ("SEULA") contains additional terms and conditions for the Software product licensed under the End User License Agreement ("EULA") between You ("You" as used herein means You and the business entity you represent or "Company") and Cisco (collectively, the "Agreement"). Capitalized terms used in this SEULA but not defined will have the meanings assigned to them in the EULA. To the extent that there is a conflict between the terms and conditions of the EULA and this SEULA, the terms and conditions of this SEULA will take precedence.

In addition to the limitations set forth in the EULA on your access and use of the Software, you agree to comply at all times with the terms and conditions provided in this SEULA.

DOWNLOADING, INSTALLING, OR USING THE SOFTWARE CONSTITUTES ACCEPTANCE OF THE AGREEMENT, AND YOU ARE BINDING YOURSELF AND THE BUSINESS ENTITY THAT YOU REPRESENT TO THE AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THE AGREEMENT, THEN CISCO IS UNWILLING TO LICENSE THE SOFTWARE TO YOU AND (A) YOU MAY NOT DOWNLOAD, INSTALL OR USE THE SOFTWARE, AND (B) YOU MAY RETURN THE SOFTWARE (INCLUDING ANY UNOPENED CD PACKAGE AND ANY WRITTEN MATERIALS) FOR A FULL REFUND, OR, IF THE SOFTWARE AND WRITTEN MATERIALS ARE SUPPLIED AS PART OF ANOTHER PRODUCT, YOU MAY RETURN THE ENTIRE PRODUCT FOR A FULL REFUND. YOUR RIGHT TO RETURN AND REFUND EXPIRES 30 DAYS AFTER PURCHASE FROM CISCO OR AN AUTHORIZED CISCO RESELLER, AND APPLIES ONLY IF YOU ARE THE ORIGINAL END USER PURCHASER.

For purposes of this SEULA, the Product name and the Product description You have ordered is any of the following Cisco Systems Email Security Appliance ("ESA"), Cisco Systems Web Security Appliance ("WSA") and Cisco Systems Security Management Application ("SMA") (collectively, "Content Security") and their Virtual Appliance equivalent ("Software"):

Cisco AsyncOS for Email

Cisco AsyncOS for Web

Cisco AsyncOS for Management

Cisco Email Anti-Spam, Sophos Anti-Virus

Cisco Email Outbreak Filters

Cloudmark Anti-Spam

Cisco Image Analyzer

McAfee Anti-Virus

Cisco Intelligent Multi-Scan

Cisco RSA Data Loss Prevention

Cisco Email Encryption

Cisco Email Delivery Mode

Cisco Web Usage Controls

Cisco Web Reputation

Sophos Anti-Malware

Webroot Anti-Malware

McAfee Anti-Malware

Cisco Email Reporting

Cisco Email Message Tracking

Cisco Email Centralized Quarantine

Cisco Web Reporting

Cisco Web Policy and Configuration Management

Cisco Advanced Web Security Management with Splunk

Email Encryption for Encryption Appliances

Email Encryption for System Generated Bulk Email

Email Encryption and Public Key Encryption for Encryption Appliances

Large Attachment Handling for Encryption Appliances

Secure Mailbox License for Encryption Appliances

Definitions

For purposes of this SEULA, the following definitions apply:

"Company Service" means the Company's email, Internet, security management services provided to End Users for the purposes of conducting Company's internal business.

"End User" means: (1) for the WSA and SMA, the employee, contractor or other agent authorized by Company to access the Internet and the SMA via the Company Service; and (2) for the ESA, the email boxes of the employees, contractors, or other agent authorized by Company to access or use the email services via the Company Service.

"Ordering Document" means the purchase agreement, evaluation agreement, beta, pre-release agreement or similar agreement between the Company and Cisco or the Company and a Cisco reseller, or the valid terms of any purchase order accepted by Cisco in connection therewith, containing the purchase terms for the Software license granted by this Agreement.

"Personally Identifiable Information" means any information that can be used to identify an individual, including, but not limited to, an individual's name, user name, email address and any other personally identifiable information.

"Server" means a single physical computer or devices on a network that manages or provides network resources for multiple users.

"Services" means Cisco Software Subscription Services.

"Service Description" means the description of the Software Subscription Support Services at http://www.cisco.com/web/about/doing_business/legal/service_descriptions/index.html

"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

Symbols

/dev/null, in alias tables 11-2

A

Access Policies 8-8
active sessions 12-21
administration
  delegating 12-4
administration commands 13-1
admin password
  changing 2-10
Advanced File Publishing
  when to use 8-1
alertlisting 13-33
alert messages 2-9
alert recipient 13-28
alerts
  alert classifications 13-29
  recipients 13-28
  settings 13-28
  severities 13-29
alert settings 2-9
alternate MX host 11-1
AMW
  See anti-malware
anonymize user names 5-4
anti-malware 5-31
appliance status
  Web Security appliances 8-19
archiving reports 4-56, 4-58, 5-62, 5-66
AsyncOS reversion 13-25
AsyncOS update servers 13-20
AsyncOS upgrades 13-24
automatic update
  interval 13-20
automatic updates 13-20
AutoSupport feature 2-10, 13-30

B

backing up 13-6
backup
  scheduling 13-9
backupconfig command 13-49
browser
  multiple windows or tabs 2-7
browser requirements 2-6
Bypass Settings 8-8

C

case-sensitivity
  in LDAP queries 10-8
Centralized Configuration Management 8-1
chain query
  creating 10-10
  LDAP 10-10
Change Password link 12-12
changing your password 12-12
clean message 4-13, 4-16
CLI Audit Logs 14-5
Client Malware Risk 5-37
Client Malware Risk report 5-37
Client Malware Risk report page 5-37

comments 11-5
comments in imported files 11-5
configuration
overview 2-1
publishing to Web Security appliances 8-13
configuration file 13-43
CLI 13-48
XML 13-44
Configuration Master
assigning Web Security appliances to 8-5
associating Web Security appliances 8-5
configuring Web Security features 8-8
preconfiguring 8-6
publishing 8-13
when to use 8-1
Configuration Master 7.1 8-8
Configuration Master 7.5 8-8
Configuration Master 7.7 8-8
configuring the appliance
reporting 5-1
scheduling reports 4-55, 5-62
Custom URL Categories 8-8
custom URL categories
overview 5-24

D

daily magnitude 4-20
Data Security 8-8
Decryption Policies 8-8
default
gateway 2-11
hostname 2-11
IP address 2-8
router 2-11
default DNS server 13-38
Defined Time Ranges 8-8
delegated administration 12-4
delegating administration 12-4, 12-7
delivery 11-1
disaster recovery 13-12
disk quotas
editing 13-50
DLP Incident Summary page 4-30
DNS C-1
authoritative server 13-37
disabling reverse DNS lookup timeout 13-38
double lookup 4-19
priority 13-37
servers 2-11
setting 2-11
splitting 13-37
timeout 13-37
timeout for reverse DNS lookups 13-38
DNS cache, flushing 13-38
dnsconfig command 13-37
dnsflush command 13-38
DNS servers 13-37
DNS settings 13-39
Domain-Based Executive Summary Report 4-52
Domain Name Service (DNS)
settings 2-11
Domain Redirect feature, see smtproutes command
domains 4-20
double-DNS verified 4-18
DTD (document type definition) 13-46

E

editing disk quotas 13-50
editing DNS settings via GUI 13-39
edit packet capture 15-4
email
clean message 4-13, 4-16
email reporting groups 4-4
Email Security Appliance
adding as managed appliance 4-3, 6-3, 7-6
Email Security Monitor

Items Displayed menu 4-18
Envelope Recipient 6-5
Envelope Sender 6-4
Ethernet interfaces B-1
event tracking 6-5
Currently in Outbreak Quarantine 6-5
Delivered 6-5
DLP Violations 6-5
Hard Bounced 6-5
Quarantined as Spam 6-5
Soft Bounced 6-5
Spam Positive 6-5
Suspect Spam 6-5
Virus Positive 6-5
exporting
reports 3-8, 3-9
external authentication 10-14
enabling LDAP 12-16
enabling RADIUS 12-16
External Data Loss Prevention 8-8
External DLP Policies 8-8

F

feature keys 8-19, 13-2
adding (GUI) manually 13-2
firewall 2-5
firewall ports C-1
FTP C-1
FTP Access A-3
FTP Poll 14-2
FTP Push 14-2
FTP Server Logs 14-5

G

globbing 11-2
GUI
browser requirements 2-6
GUI logs 14-5

H

HAT
Host Access Table
hostname, setting 13-36
HTTP C-1
HTTP Logs 14-5
HTTP proxy server 13-20
HTTPS proxy server 13-21

I

Identities 8-8, 8-9, 8-14
IMAP authentication 7-9
installation
reverting 13-25
invalid recipient 4-13, 4-15
IP address profile pages 4-19
IPv6 4-5, 6-4, 6-5, 7-15

K

keys 13-2

L

L4TM 8-2
L4 Traffic Monitor
report 5-42
language
preference 13-52
specifying a default language for Cisco IronPort Spam Quarantine 7-3
last command 12-22
LDAP C-2
alias consolidation queries 10-6

chain queries 10-10
domain-based queries 10-8
end-user authentication queries 10-5
external authentication 10-14, 12-16
failover 10-11
LDAP server profiles 10-2
load-balancing 10-11
multiple servers 10-11
overview 10-1
testing queries 10-8
test servers 10-4
LDAPS C-2
Global Catalog Server C-2
Level 4 Traffic Monitor 8-2
limits
SMTP Routes 11-3
loadconfig command 13-49
log file type 14-4
logging
overview 14-1
logging versus reporting 14-1
logheaders command 14-25
logs
Cisco IronPort Spam Quarantine GUI Logs 14-6
Cisco IronPort Spam Quarantine Logs 14-6
Cisco IronPort Text Mail Logs 14-5
CLI Audit Logs 14-5
comparison 14-7
Configuration History Logs 14-7
definition 14-1
extensions in filenames 14-26
format 14-1
FTP Server Logs 14-5
global attributes 14-24
HTTP Logs 14-5
Injection Debug Logs 14-6
levels 14-23
log subscription defined 14-4
message headers in 14-25
NTP Logs 14-5
Reporting Logs 14-5
Reporting Query Logs 14-6
Safelist/Blocklist Logs 14-6
SCP Push 14-2
SMA Logs 14-6
Status Logs 14-6
subscriptions 14-2
syslog push 14-2
log subscriptions 14-2, 14-4

M

mailconfig command 13-48
mailertable feature 11-1
MAIL FROM
configuring for notifications 13-28
mail trend graph 4-11
malware
types blocked 5-35
mapping domains 11-1
McAfee
update servers 13-20
message headers 14-25
message tracking
See tracking
message variables
Cisco IronPort Spam Quarantine notifications 7-10
monitoring
scheduling reports 4-55, 5-62
summary data 4-1, 5-1
monitoring services
enabling on Security Management Appliance 2-14
M-Series appliance 2-3

N

netmasks, selecting B-1

network_access_list 12-18
networking worksheet 2-5
network owner 4-20
Network Owner profile pages 4-19
network time protocol (NTP)
  settings 2-10
network topology B-4
No Subject 6-8
NTP C-1
NTP Logs 14-5
NTP server 13-43

O

offline command 13-4
offline state 13-4
on-demand reports 5-65
Outbound Malware Scanning 8-8
Outbreak Heuristics 5-36
Outgoing Destinations page 4-24
Outgoing Senders page 4-26
Overall Bandwidth Limits 8-8
Overview page
  Email reporting 4-10, 4-14
  Web reporting 5-12

P

packet capture 15-3
password
  changing 12-12
password command 13-40
passwords
  length 12-10
passwords, changing 13-40
policy groups
  custom URL categories 5-24
POP authentication 7-9
power down 13-3
preferences
  setting 13-52
Profile for Domain pages 4-19
Proxy Buffer Memory 5-59
Proxy Bypass 8-8
proxy server 13-20
publishconfig command 13-49
publish history
  viewing 8-18
publishing configurations
  advanced file publish 8-16
  Configuration Master 8-13
  to Web Security appliances 8-13
  viewing history 8-18

Q

quarantine 7-1
queries
  chain queries 10-10
  domain-based 10-8
  external authentication 10-14
  LDAP alias consolidation 10-6
  LDAP end-user authentication 10-5

R

RADIUS external authentication 12-16
RAT
  Recipient Access Table
reboot command 13-3
reconfigure 2-8
recursive DNS queries 13-38
recursive entries
  in SMTP Routes 11-2
redirecting email 11-1
remote upgrades 13-17

removing 13-43
Reporting Logs 14-5
Reporting Query Logs 14-6
Reports
print 3-8
reports
archiving 4-56, 4-58, 5-62, 5-66
charts 3-5
Client Malware Risk 5-37
Client Malware Risk Page 5-37
csv 3-8, 3-9
exporting data 3-8, 3-9
filters 3-7
graphs 3-5
interactive display 5-1
interactive pages
time range 3-4
L4 Traffic Monitor 5-42
Malware Category 5-33
Malware Threat 5-34
on-demand 5-65
pdf 3-8
performance 3-7
printing 3-8
scheduling 4-55, 5-62
time range
preferences 13-52
time range for scheduled reports 4-55, 5-62
uncategorized URLs 5-26
URL Categories 5-24
Web Reputation Filters 5-39
resetconfig command 13-5
resetting 13-5
resume command 13-4
reverse DNS 6-8
Reverse DNS Lookup
disabling 13-38
reverse DNS lookup
timeout 13-37
revert
installation 13-25
rolling over log files 14-2
rollovernow command 14-26
root servers (DNS) 2-11
routing 11-1
Routing Policies 8-8
routing taking precedence over selected interface B-3

S

SaaS Policies 8-8, 8-14
Safelist/Blocklist Logs 14-6
saveconfig command 13-48
SBRS score 6-8
scheduling a backup 13-9
instant 13-10
interruption of backup process 13-8
periodic 13-9
scp command A-6
SCP Push 14-2
secure copy A-6
Security Management Appliance
backing up 13-6
enabling services 2-14
Security Management appliance
Incoming Mail Graph 4-12
Incoming Mail Summary 4-12
Outgoing Mail Graph 4-12
Outgoing Mail Summary 4-12
Security Services Display page 8-11
security services settings
editing 8-11
SenderBase 4-14, 4-19, 4-20, 6-8, C-1
sender groups 4-23
serial connection pinouts A-7
services for interfaces A-1
sethostname command 13-36
showconfig command 13-48

shutdown command 13-3
shutting down 13-3
SMA Logs 14-6
SMTP C-1
SMTP Authentication 6-8
SMTP Routes 11-1
  deleting all 11-4
  limits 11-3
  mail delivery and splintering 11-3
  multiple host entries 11-2
  recursive entries in 11-2
  USEDNS 11-3
SMTP Routes, maximum 11-1
SMTP Routes and DNS 11-3
spam message 4-13, 4-15
Spam Quarantine, Cisco IronPort
  default language 7-3
  defined 7-1
  deleting all messages 7-18
  end user access without authentication 7-9
  end-user authentication queries 10-5
  GUI Logs 14-6
  Logs 14-6
  message details 7-18
  message variables 7-10
  notification 7-1
  released messages and email pipeline 7-18
specifying an offset 13-41
SSH C-1
start packet capture 15-3
Status Logs 14-6
stopped by content filter 4-9, 4-13, 4-16
stopped by reputation filtering 4-13, 4-15
streaming upgrades 13-17
subject
  no subject 6-8
supportrequest command 15-1
suspend command 13-4
synchronizing time 2-10
Syslog 14-2
system administration 13-1
System Capacity
  All page 4-49
  Incoming Mail page 4-45
  memory page swapping 4-48
  Outgoing Mail page 4-46
  System Load page 4-47
  WorkQueue page 4-44
System Capacity page
  for ESA 4-43
system clock 2-10
system failure
  disaster recovery on Security Management Appliance 13-12
System Logs 14-6
system time
  setting 2-10

T

tail command 14-27
  parameters 14-27
Telnet C-1
Text Mail Logs, Cisco IronPort 14-5
tiered reporting 4-4
time, system 2-10
Time Keeping Method 13-42
Time range
  configure 3-4
Time Ranges 8-8
time servers 2-10
time zone 13-41
time zone, setting 2-10
time zone file updates
  overview 13-42
Time Zone page 13-41
TLS Connections page 4-7, 4-36
tracking

advanced options 6-4
event 6-5
message details 6-4
result set, narrowing 6-6
Transparent User Identification 8-14
turning off 13-3

U

uncategorized URLs
in reports 5-26
update server 13-20
Update Settings page 13-20
updating
time zone files 13-42
upgrades C-1
remote 13-17
streaming 13-17
upgrade server 13-17
URL categories
uncategorized URLs 5-26
URL Categories report 5-24
URL category set
updates 8-22, 13-28
URL Filters
custom categories 5-24
user accounts 12-9, 12-10, 12-16
locking and unlocking 12-13, 12-15
user groups 12-2
user name 12-10
user names
anonymous 5-4
user roles 12-2
custom 12-4
custom, for Email 12-4
custom, for Web 12-7
description 12-2
UTF-8 6-5

V

virus message 4-13, 4-16
Virus Types page 4-34

W

WBRS (Web-Based Reputation Score) 5-53
Web Reporting
Overview page 4-10, 5-12
Web Reputation Filters
report 5-39
Web Security Appliance
adding as managed appliance 5-3, 8-5
process for managing 8-2
publishing configurations to 8-13
viewing status 8-19
Web UI session timeout 12-20
whoami command 12-22
who command 12-22

X

XML 13-43, 13-44, 13-46, 13-48