IRONPORT®
SYSTEMS

IronPort AsyncOS™ 5.6.2

USER GUIDE

for Web Security Appliances

COPYRIGHT

Part Number: 421-0527

Revision Date: December 22, 2008

IRONPORT SYSTEMS® , INC.

IronPort Systems, Inc.
950 Elm Ave.
San Bruno, CA 94066

CONTACTING IRONPORT CUSTOMER SUPPORT

If you have purchased support directly from IronPort Systems, you can request our support by phone, email or online 24 hours a day, 7 days a week. During our office hours (24 hours per day, Monday through Friday excluding US holidays), one of our engineers will contact you within an hour of your request. To report a critical issue that requires urgent assistance outside of our office hours, please call us immediately at the numbers below.

U.S. Toll-free: 1 (877) 641-IRON (4766)

International: `www.ironport.com/support/contact_support.html`

Support Portal: `www.ironport.com/support`

If you have purchased support through a reseller or another entity, please contact them for support of your IronPort products.

Table of Contents

List of Figures

List of Tables

CHAPTER 1

Getting Started with the Web Security Appliance

The *IronPort AsyncOS for Web User Guide* provides instructions for setting up, administering, and monitoring the IronPort Web Security appliance. These instructions are designed for an experienced system administrator with knowledge of networking and web administration.

This chapter discusses the following topics:

WHAT'S NEW IN THIS RELEASE

This section describes the new features and enhancements in AsyncOS for Web 5.6.2. For more information about the release, see the product release notes, which are available on the IronPort Customer Support Portal at the following URL:

`http://www.ironport.com/support/login.html`

Note — You need a Support Portal account to access the site. If you do not already have an account, click the Request an Account link on the Support Portal login page. Generally, only IronPort customers, partners, and employees can access the Support Portal.

You might also find it useful to review release notes for earlier releases to see the features and enhancements that were previously added. To view those release notes on the Support Portal, click the Earlier Releases link on the appropriate appliance documentation page.

For a description of new features and enhancements in AsyncOS for Web 5.6.0 and 5.6.1, see "What's New in Version 5.6.0" on page 5 and "What's New in Version 5.6.1" on page 3.

New Feature: Enable Web Proxy on P2

In AsyncOS 5.6.2 for Web, you can enable the Web Proxy so it listens for web requests on the P2 network interface. By default, the Web Proxy does not listen for requests on P2, even when enabled. However, you can configure it to listen for requests on P2 using the `advancedproxyconfig > miscellaneous` CLI command.

For more information, see "Miscellaneous Options" on page 80.

WHAT'S NEW IN VERSION 5.6.1

This section describes new features and enhancements added in the AsyncOS 5.6.1 for Web release.

New Feature: Secure Client Authentication

In AsyncOS 5.6.1 for Web, clients can send authentication credentials to the Web Security appliance securely, even when using Basic authentication scheme.

When NTLMSSP authentication is used, authentication credentials are always passed to the Web Proxy securely. When Basic authentication is used, authentication credentials are passed to the Web Proxy insecurely by default. They are encoded, but not encrypted, and sent over HTTP. However, you can configure the Web Security appliance so clients send authentication credentials securely. This works for both LDAP and NTLM Basic authentication.

For more information, see "Sending Authentication Credentials Securely" on page 279.

New Feature: External Authentication Support

In AsyncOS 5.6.1 for Web, you can configure the Web Security appliance to use an external RADIUS directory to authenticate users logging in to the appliance. You can configure the appliance to contact multiple external servers for authentication. You might want to define multiple external servers to allow for failover in case one server is temporarily unavailable. When you define multiple external servers, the appliance connects to the servers in the order defined on the appliance.

For more information, see "Using External Authentication" on page 374.

Enhanced: PAC File Support

In AsyncOS 5.6.1 for Web, when browsers are configured to use a PAC file on the appliance, the URL should include the PAC file name. If the URL does not specify the PAC file name, the appliance uses default.pac if it exists and returns an error if it does not.

Enhanced: Cookie Based Authentication in Explicit Forward Mode

In AsyncOS 5.6.1 for Web, you can configure the Web Proxy to use cookie based authentication when it is deployed in explicit forward mode. In previous versions, cookie based authentication was available only when the Web Proxy was deployed in transparent mode.

To do this, use the `advancedproxyconfig > authentication` CLI command and enable the following option:

```
Would you like to use surrogates for explicit forward mode requests?
```

After you enable surrogates, you can configure them in the rest of the `advancedproxyconfig > authentication` CLI command.

For more information, see "Authentication Options" on page 72.

Enhanced: Intermediate Digital Certificate Support for HTTPS Decryption

In AsyncOS 5.6.1 for Web, you can upload an intermediate certificate that has been signed by a root certificate authority in addition to loading a root certificate. When the Web Proxy mimics the server certificate to decrypt an HTTPS transaction, it sends the uploaded certificate along with the mimicked certificate to the client application. That way, as long as the intermediate certificate is signed by a root certificate authority that the client application trusts, the application will trust the mimicked server certificate, too. You might want to upload an intermediate certificate if your organization uses its own root certificate authority, but does not want to upload the root certificate to the Web Security appliance for security reasons.

WHAT'S NEW IN VERSION 5.6.0

This section describes new features and enhancements added in the AsyncOS 5.6.0 for Web release.

Enhanced: Policy Groups

In AsyncOS 5.6 for Web, you have greater flexibility in defining policies that control access to the web and enforce organizational policies and requirements. AsyncOS 5.6 for Web includes additional types of policies, and the existing policies have changed names. The policy types include:

- **Identities.** An identity is a policy that identifies and groups users. It addresses the question, "who are you?"
- **Access policies.** These policies are renamed from web access policies.
- **Decryption policies.** These policies are renamed from HTTPS decryption policies. It addresses the question, "to decrypt or not to decrypt?"
- **Routing policies.** A routing policy determines to where to pass the client request, either to another proxy or to the destination server. It addresses the question, "from where to fetch content?"

For more information about working with policies, see "Working with Policies" on page 83.

In addition to the new policy types, you can also define policy group membership by additional criteria.

Time Based Policies

Time-based policies are a valuable feature for acceptable use policy enforcement. They enable customers to configure different URL filtering actions for different times of the day (or days of the week). For more information about time based policies, see "Working with Time Based Policies" on page 93.

User Agent Based Policies

Some applications behave abnormally when the Web Security appliance applies certain policies (e.g. authentication). Customers can now configure policies (e.g. bypass authentication) specific to these applications, which are identified by their user agent strings.

New Feature: Policy Trace Tool

In AsyncOS 5.6 for Web, customers with many policies configured in the Web Security Manager can use the new policy trace tool to simulate a transaction going through the Web Security appliance. By entering client and destination information for a transaction, the policy trace tool tells the customer what policies matched, what policies did not, and ultimately, what policies were applied. Customers now have an easy way to troubleshoot and debug their policies on the Web Security appliance.

For more information about the policy trace tool, see "Tracing Policies" on page 98.

New Feature: Custom IronPort Notification Pages

Some customers may want to customize the end-user notification pages that are displayed to end users, but do not want to deploy a separate web server to host those pages. In AsyncOS 5.6 for Web, the Web Security appliance can now host basic HTML pages that you create (no JavaScript, no external objects) locally. These custom pages override the default IronPort pages shipped with the Web Security appliance.

For more information about customizing IronPort notification pages, see "Editing IronPort Notification Pages" on page 187.

New Feature: Proxy Bypass List

In AsyncOS 5.6 for Web, you can create a list of destination addresses that bypass all Web Proxy security and access policy features, such as web reputation, URL filtering, and anti-malware scanning. The proxy bypass list provides customers for a quick workaround to the destination that does not consistently work well with a web proxy that is deployed in transparent mode.

For more information about creating a proxy bypass list, see "Bypassing the Web Proxy" on page 61.

New Feature: Auto Feature Key Download

In AsyncOS 5.6 for Web, the Web Security appliance periodically queries the IronPort servers for new or refreshed feature keys. If one exists, the appliance automatically downloads and applies the feature key without any administrative intervention.

For more information about working with feature keys, see "Working with Feature Keys" on page 369.

Enhanced: SNMP Support

In AsyncOS 5.6 for Web, the Web Security appliance now supports SNMPv1, v2, and v3. Specifically, the following capabilities have been introduced to offer an enterprise-grade monitoring system:

- New MIB objects for hardware and software monitoring.
- SNMP traps for hardware and software alerts.
- Community strings to enforce read-only access control over MIB objects.

For more information about using SNMP, see "SNMP Monitoring" on page 317.

Enhanced: Multiple Upstream Proxy Support

In AsyncOS 5.6 for Web, you can configure multiple upstream proxies for failover, load balancing, and conditional routing (specific clients or destinations routed to a specific upstream proxy). This feature enables the Web Security appliance to deploy easily in an environment with a tiered proxy hierarchy.

For more information about working with multiple upstream proxies, see "Routing Traffic to Upstream Proxies" on page 139.

Enhanced: Active Mode FTP over HTTP

AsyncOS 5.6 for Web supports active mode FTP over HTTP for those organizations that allow active mode FTP connections. (Note: Passive mode FTP over HTTP is already supported.) To enable active FTP, use the `advancedproxyconfig > ftp` CLI command.

For more information about the `advancedproxyconfig` command options, see "Advanced Proxy Configuration" on page 71.

Enhanced: Logging

AsyncOS 5.6 for Web includes several changes and enhancements to Web Security appliance logging to help you troubleshoot issues more easily.

Web Proxy Logging

AsyncOS 5.6 for Web includes multiple new log file types that each record information for a different aspect of the Web Proxy. This allows you to turn on logging for the specific Web Proxy module that is showing signs of an error. Any "Proxy Log" log subscription from a previous version gets upgraded to the "Default Proxy Logs."

If a user or administrator encounters an issue with the Web Proxy behavior, read the Default Proxy Logs first. If you see a log entry that you suspect might be the symptom of an issue, then you can create a log subscription for the relevant specific Web Proxy module. Then read that proxy log to help troubleshoot the problem.

For more information, see "Web Proxy Logging" on page 334.

New Log File Types

AsyncOS 5.6 for Web includes the following new types of log files:

- **Authentication Framework Logs.** Records authentication history and messages.
- **NTP Logs.** Records changes to the system time made by the Network Time Protocol.
- **Proxy Bypass Logs.** Records transactions that bypass the Web Proxy.
- **Status Logs.** Records information related to the system, such as feature key downloads.

For more information, see "Log File Types" on page 332.

Enhanced: Better Troubleshooting and Diagnostics Tools

In AsyncOS 5.6 for Web, better tools have been implemented for customers and IronPort Customer Support to troubleshoot common deployment issues, such as authentication and networking. These tools are designed to reduce ticket resolution time, and thus, get you up and running faster.

HOW TO USE THIS GUIDE

Use this guide as a resource to learn about the features of your IronPort appliance. The topics are organized in a logical order. You might not need to read every chapter in the book.

You can also use this guide as a reference book. It contains important information, such as network and firewall configuration settings, that you can refer to throughout the life of the appliance.

The guide is distributed in print and electronically as PDF and HTML files. The electronic versions of the guide are available on the IronPort Customer Support Portal. You can also access the HTML online help version of the book directly from the appliance GUI by clicking the Help and Support link in the upper-right corner.

Before You Begin

Before you read this guide, review the *IronPort Quickstart Guide* and the latest product release notes for your appliance. In this guide, it is assumed that you have unpacked the appliance, physically installed it in a rack, and turned it on.

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

Typographic Conventions

Typeface	Meaning	Examples
`AaBbCc123`	The names of commands, files, and directories; on-screen computer output.	`Please choose an IP interface for this Listener.` The `sethostname` command sets the name of the IronPort appliance.

`AaBbCc123`	User input, in contrast to on-screen computer output.	`mail3.example.com> commit`
		`Please enter some comments describing your changes:`
		`[]> Changed the system hostname`

AaBbCc123	Book titles, new terms, emphasized words, and command line variables; for command line variables, the italicized text is a placeholder for the actual name or value.	Read the *IronPort Quickstart Guide*.
		
		The IronPort appliance *must* be able to uniquely select an interface to send an outgoing packet.
		
		`Before you begin, please reset your password to a new value.`
		`Old password: ironport`
		`New password: your_new_password`
		`Retype new password: your_new_password`

Where to Find More Information

IronPort offers the following resources to learn more about the Web Security appliance.

Documentation Set

The documentation for the Web Security appliance includes the following books:

- *IronPort AsyncOS for Web User Guide* (this book)
- *IronPort AsyncOS CLI Reference Guide*

Occasionally, this book refers to the other guides for additional information about topics.

Knowledge Base

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

`http://www.ironport.com/support/login.html`

Note — You need a Support Portal account to access the site. If you do not already have an account, click the Request an Account link on the Support Portal login page. Generally, only IronPort customers, partners, and employees can access the Support Portal.

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

Articles generally fall into one of the following categories:

- **How-To.** These articles explain how to do something with an 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 an 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 IronPort products. For example, a troubleshooting article might provide steps to follow if you are having problems with DNS.

Each article in the Knowledge Base has a unique answer ID number.

IronPort Nation

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

You access IronPort Nation on the Customer Support Portal at the following URL:

`http://www.ironport.com/support/login.html`

IronPort Customer Support

You can request IronPort product support by phone, email, or online 24 hours a day, 7 days a week.

During Customer Support hours — 24 hours a day, Monday through Friday, excluding U.S. holidays — an engineer will contact you within an hour of your request.

To report a critical issue that requires urgent assistance outside of Customer Support hours, contact IronPort using one of the following methods:

U.S. Toll-free: 1 (877) 641-IRON (4766)

International: `http://www.ironport.com/support/contact_support.html`

Support Portal: `http://www.ironport.com/support/login.html`

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

IronPort Welcomes Your Comments

The 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:

`docfeedback@ironport.com`

Please include the following part number in the subject of your message: 421-0527.

WEB SECURITY APPLIANCE OVERVIEW

The Web Security appliance is a robust, secure, efficient device that protects corporate networks against web-based malware and spyware programs that can compromise corporate security and expose intellectual property. The Web Security appliance extends IronPort's SMTP security applications to include protection for standard communication protocols, such as HTTP, HTTPS, and FTP.

Malware ("malicious software") is software designed to infiltrate or damage a computer system without the owner's consent. It can be any kind of hostile, intrusive, or annoying software or program code. Web-based malware includes spyware, system monitors, adware, phishing and pharming techniques, keystroke (key) loggers, browser hijackers, trojan horses, and more.

Web-based malware is a rapidly growing threat, responsible for significant corporate downtime, productivity losses and major strains on IT resources. Additionally, companies run the risk of violating compliance and data privacy regulations if their networks become compromised by malware. Companies run the risk of expensive legal costs and exposure of intellectual property.

The best place to stop these threats from entering the network is right at the gateway. The Web Security appliance provides deep application content inspection, by offering a web proxy service and by monitoring layer 4 traffic. The Web Proxy and Layer 4 Traffic Monitor allow organizations to ensure breadth of coverage within their networks. The Web Security appliance provides a powerful web security platform to protect your organization against malware that is optimized for performance and efficacy.

CHAPTER 1

Using the Web Security Appliance

This chapter contains the following topics:

HOW THE WEB SECURITY APPLIANCE WORKS

The Web Proxy and the L4 Traffic Monitor are independent services. They are enabled and configured separately to provide the highest level of protection against a broad range of web-based malware threats.

The Web Proxy and L4 Traffic Monitor use data that is stored in filtering tables to evaluate and match URL request attributes such as domain names, and IP address path segments with locally maintained database records. If a match occurs, access policy settings determine an action to block or monitor the traffic. If no match occurs, processing continues.

Web Proxy

The Web Security appliance Web Proxy supports the following security features:

- Policy groups — Policy groups allow administrators to create groups of users and apply different levels of category-based access control to each group.
- IronPort URL Filtering Categories — IronPort URL Filters allow you to configure how the appliance handles each web transaction based on the URL category of a particular HTTP request.
- Web Reputation Filters — Reputation filters analyze web server behavior and characteristics to identify suspicious activity and protect against URL-based malware threats.
- Anti-Malware Services — The IronPort DVS™ engine in combination with the Webroot™ and McAfee scanning engines identify and stop a broad range of web-based malware threats.

For detailed information about Web Proxy services, see “Web Proxy Services” on page 55.

The L4 Traffic Monitor

The L4 Traffic Monitor is a configurable service that listens and monitors network ports for rogue activity and blocks malware attempts to infect your corporate network. Additionally, the L4 Traffic Monitor detects infected clients and stops malicious activity from going outside the corporate network.

For detailed information about the L4 Traffic Monitor, see “L4 Traffic Monitor” on page 295.

ADMINISTERING THE WEB SECURITY APPLIANCE

You can manage the Web Security appliance using a web-based administration tool. When you first access the appliance, the web interface launches the System Setup Wizard to perform an initial configuration. After running the System Setup Wizard, you can use the web interface or Command Line Interface (CLI) to customize settings and maintain your configuration.

For a description of how to access the CLI and a list CLI supported commands, see "Command Line Interface" on page 415.

System Setup Wizard

The System Setup Wizard is a utility that configures basic settings and enables a set of system defaults. The System Setup Wizard is located on the System Administration tab. For more information about running the System Setup Wizard, see "System Setup Wizard" on page 37.

Note — Running the System Setup Wizard completely reconfigures the Web Security appliance and resets the administrator password. Only use the System Setup Wizard the first time you install the appliance, or if you want to completely overwrite the existing configuration. Running the System Setup Wizard after the appliance is already configured can also interrupt client access to the web. If you choose to run the System Setup Wizard after performing an initial setup, use the System Administration > Configuration File pages to print a configuration summary and archive the current configuration file.

Accessing the Web Security Appliance

To access the appliance and launch the web-based administration utility, open a web browser. For the list of supported web browsers, see "Browser Requirements" on page 7.

Connect to the management interface using one of the following methods:

- IP address and port number

  `http://192.168.42.42:8080`

  where `192.168.42.42` is the default IP address, and `8080` is the default admin port setting.

- Host name and port number

  `http://hostname:8080`

  where `hostname` is the name of the appliance, and `8080` is the default admin port setting.

Note — The hostname parameter is assigned during system setup. Before you can connect to the management interface using a hostname, you must add the appliance hostname and IP address to your DNS server database.

For information about how to use and navigate the web interface, see "Navigating the Web Security Appliance Web Interface" on page 5.

Using the Command Line Interface (CLI)

You can establish an SSH or serial console connection to administer the appliance using the CLI. The Web Security appliance CLI supports a set of commands to access, install, and administer the system. See "Command Line Interface" on page 415 for information about the CLI and a list of supported commands that can be used to access, upgrade, and administer the appliance.

The SenderBase Network

The SenderBase Network is a threat management database that tracks millions of domains around the world and maintains a global watch list for Internet traffic. SenderBase provides IronPort with an assessment of reliability for known Internet domains. The Web Security appliance uses the SenderBase data feeds to improve the accuracy of Web Reputation Scores.

Basic SenderBase Network Participation is enabled by default during system setup. The appliance supports three levels of participation in the SenderBase Network:

- **Disabled.** Participation is disabled and none of the data that the appliance collects is sent back to the SenderBase Network servers.
- **Limited.** Basic participation summarizes server name information and sends MD5-hashed path segments to the SenderBase Network servers.
- **Standard.** Enhanced participation sends the entire URL with unobfuscated path segments to the SenderBase Network servers. This option assists in providing a more robust database, and continually improves the integrity of Web Reputation Scores.

To select a level of participation in the SenderBase Network, use the Security Services > SenderBase page.

Sharing Data

Participating in the SenderBase Network means that IronPort collects data and shares that information with the SenderBase threat management database. This data includes information about request attributes and how the appliance handles requests.

IronPort recognizes the importance of maintaining your privacy, and does not collect or use personal or confidential information such as usernames and passwords. Additionally, the file names and URL attributes that follow the hostname are obfuscated to ensure confidentiality. When it comes to decrypted HTTPS transactions, the SenderBase Network only receives the IP address, web reputation score, and URL category of the server name in the certificate.

If you agree to participate in the SenderBase Network, data sent from your IronPort appliance is transferred securely using HTTPS. Sharing data improves IronPort's ability to react to web-based threats and protect your corporate environment from malicious activity.

Reporting and Logging

The Web Security appliance provides several options for capturing data and monitoring system activity. For detailed information about scheduling reports, see "Reporting Overview" on page 324. For more information about working with log files, see "Logging" on page 331.

NAVIGATING THE WEB SECURITY APPLIANCE WEB INTERFACE

The Web Security appliance web interface is a web-based administration tool that allows you to configure and monitor the appliance. The web interface allows you to configure the appliance similar to the Command Line Interface (CLI). However, some features available in the web interface are not available in the CLI and vice versa. For more information about the CLI, see "Command Line Interface" on page 415.

The Web Security appliance web interface contains multiple tabs where you can configure or monitor the appliance. You can set up access policies, schedule reports, enable features, and modify settings as necessary. The web interface also includes two menus from which you can perform basic administration tasks.

To use the web interface, open a web browser and log in. For more details, see "Accessing the Web Security Appliance" on page 3. For a list of supported web browsers, see "Browser Requirements" on page 7.

The web interface contains the following menus:

- **Options.** From this menu, you can manage your user account. You can logout or change the password you use to log in to the web interface.
- **Help.** From this menu, you can access help from documentation or IronPort Customer Support. For Help tasks, you can access the online help or the IronPort Support Portal. For Technical Support tasks, you can send a support request email to IronPort Customer Support or to allow IronPort Customer Support remote access to the Web Security appliance. For more information about the Technical Support tasks, see "Support Commands" on page 363.

The web interface contains the following tabs:

- **Monitor.** Use the pages on this tab to monitor the appliance by viewing dynamic data on website activity and appliance activity and action. For more information, see "Monitor Tab" on page 7.
- **Web Security Manager.** Use the pages on this tab to create and configure access policies that define which groups can access which types of websites. For more information, see "Web Security Manager Tab" on page 8.
- **Security Services.** Use the pages on this tab to configure how the appliance monitors and secures the network. For more information, see "Security Services Tab" on page 8.
- **Network.** Use the pages on this tab to define the network in which the appliance is located. For more information, see "Network Tab" on page 8.
- **System Administration.** Use the pages on this tab to configure administrative options, such as users, alerts, system time, and more. You can also enter keys for features you enabled during initial setup. For more information, see "System Administration Tab" on page 9.

Each tab has a list of menu selections from which you can choose. Each menu selection represents a different page in the web interface that further group information and activities. Some pages are grouped together into categories. You navigate among sections of the web interface by hovering the cursor over each tab heading and clicking a menu option from the menu that appears.

You open up other pages in the web interface by clicking on hypertext links and buttons. To find the various links, hover the cursor over text in the web interface. Links appear with an underline under the text when the cursor is over them.

Figure 1-1 on page 6 shows the web interface tabs, pages, and categories. It also shows some sample links and buttons you can click to open up other pages where you can configure the appliance.

Figure 1-1 Web Interface Tabs, Pages, and Categories

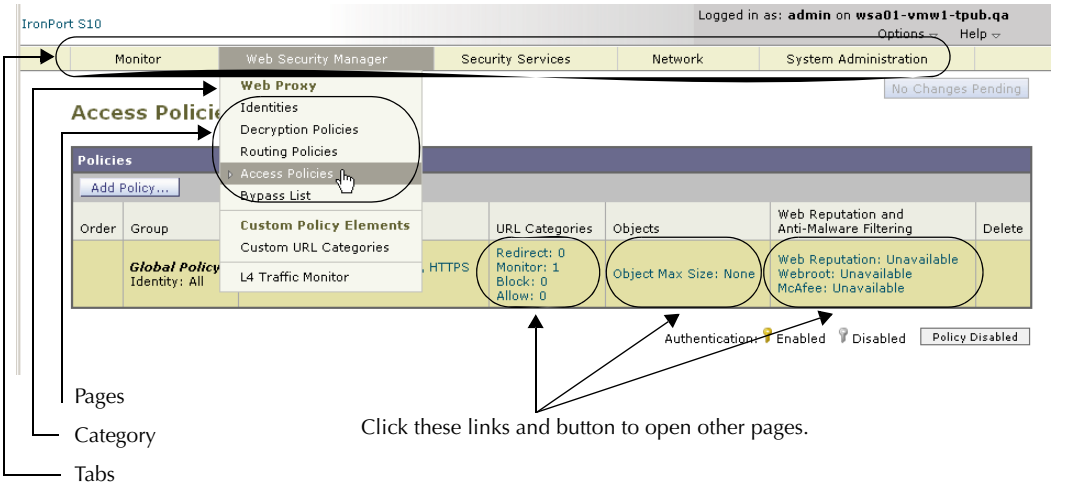

Figure 1-1 shows that the Web Security Manager tab contains the Web Proxy category, and the Web Proxy category contains the Identities, Decryption Policies, Routing Policies, Access Policies, and Bypass List pages. The tab also contains the Custom Policy Elements category (with the Custom URL Categories page), and the L4 Traffic Monitor page.

When the documentation refers to specific pages in the web interface, it uses the tab name, following by an arrow and then the page name. For example, Web Security Manager > Access Policies.

Logging In

All users accessing the web interface must log in. Type your username and password, and then click Login to access the web interface. You must use a supported web browser (see "Browser Requirements" on page 7). You can log in with the admin account or any other user

account created in the appliance. For more information creating appliance users, see "Administering User Accounts" on page 371.

After you log in, the Monitor > Overview page displays.

Browser Requirements

To access the web interface, your browser must support and be enabled to accept JavaScript and cookies. It must be able to render HTML pages containing Cascading Style Sheets (CSS). For example, you can use the following browsers:

- Firefox 1.0 and later
- Internet Explorer 6.02 and later (Windows only)
- Mozilla 1.76 and later
- Netscape 7.1 and later
- Safari 2.0.4 and later (Mac OS X only)

Your session automatically times out after 30 minutes of inactivity.

Some buttons and links in the web interface cause additional windows to open. Therefore, you may need to configure the browser's pop-up blocking settings in order to use the web interface.

Note — Only use one browser window or tab at a time to edit the appliance configuration. Also, do not edit the appliance using the web interface and the CLI at the same time. Editing the appliance from multiple places concurrently results in unexpected behavior and is not supported.

Monitor Tab

Use the Monitor tab to monitor the appliance by viewing dynamic data on website activity and appliance activity and action.

The Monitor tab includes the following pages:

- Overview
- L4 Traffic Monitor
- Web Activity
- Malware Risk
- Web Site Activity
- Anti-Malware
- URL Categories
- Web Reputation Filters
- System Status

- Report Scheduling
- Archived Reports

Web Security Manager Tab

Use the Web Security Manager tab to create and configure access policies that define which groups can access which types of websites.

The Web Security Manager tab includes the following pages:

- Identities
- Decryption Policies
- Routing Policies
- Access Policies
- Bypass List
- Custom URL Categories
- L4 Traffic Monitor Policies

Security Services Tab

Use this tab to configure how the appliance monitors and secures the network.

The Security Services tab includes the following pages:

- Proxy Settings
- HTTPS Proxy
- End-User Notification
- PAC File Hosting
- L4 Traffic Monitor
- IronPort URL Filters
- Web Reputation Filters
- Anti-Malware
- SenderBase

Network Tab

Use the Network tab to describe the network in which the appliance is located and to define the appliance's network settings.

The Network tab includes the following pages:

- Interfaces
- Transparent Redirection

- Routes
- Internal SMTP Relay
- Authentication
- Upstream Proxies
- DNS

System Administration Tab

Use the System Administration tab to configure administrative options, such as users, alerts, system time, and more. You can also enter keys for features you enabled during initial setup.

The System Administration tab includes the following pages:

- Policy Trace
- Users
- Alerts
- Log Subscriptions
- Return Addresses
- Time Zone
- Time Settings
- Configuration Summary
- Configuration File
- Feature Key Settings
- Feature Keys
- Component Updates
- Upgrade Settings
- System Upgrade
- System Setup Wizard
- Next Steps

COMMITTING AND CLEARING CHANGES

When you change the configuration of the Web Security appliance, you must commit the changes before they go into effect. Or, you can choose to clear the changes you have made if you do not want to commit them. How you commit and clear changes depends on the interface you use:

- Web interface
- Command Line Interface

Committing and Clearing Changes in the Web Interface

Commit changes using the **Commit Changes** button in the upper right corner of the web interface. You can make multiple configuration changes before you commit all of them. When you make a change, the **Commit Changes** button color is yellow and the button text changes to "Commit Changes" as shown in Figure 1-2.

Figure 1-2 The Commit Button: Changes Pending

When there are no changes to commit, the button color is gray and the button text is "No Changes Pending." Figure 1-3 shows the web interface when there are no changes to commit.

Figure 1-3 The Commit Button: No Changes Pending

You also use the **Commit Changes** button to clear the changes made since the last commit or clear.

Committing Changes

To commit changes made in the web interface:

1. Click the **Commit Changes** button.

   The Uncommitted Changes page appears.

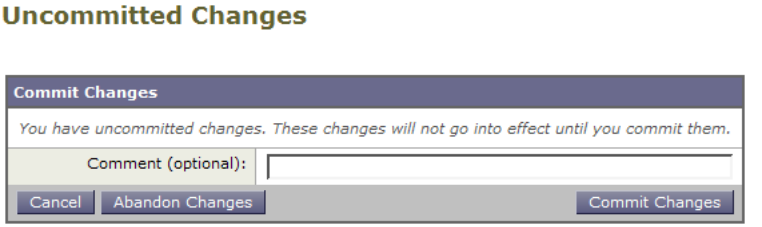

2. Enter comments in the Comment field if you choose.
3. Click **Commit Changes**.

Clearing Changes

To clear changes made in the web interface:

1. Click the **Commit Changes** button.

   The Uncommitted Changes page appears.

2. Click **Abandon Changes**.

Committing and Clearing Changes in the CLI

Commit changes using the `commit` command. Most configuration changes you make in the Command Line Interface (CLI) are not effective until you issue the `commit` command. You may include comments up to 255 characters. Changes are not verified as committed until you receive confirmation along with a timestamp. The `commit` command applies configuration changes made to appliance since the last `commit` or `clear` command issued.

For more information about using the `commit` command, see "Committing Configuration Changes" on page 420.

Clear changes using the `clear` command. For more information about using the `clear` command, see "Clearing Configuration Changes" on page 420.

CHAPTER 2

Deployment

This chapter contains the following topics:

DEPLOYMENT OVERVIEW

The Web Security appliance is typically installed as an additional layer in the network between clients and the Internet. Depending on how you deploy the appliance, you may or may not need a Layer 4 (L4) switch or a WCCP router to direct client traffic to the appliance.

When you deploy the Web Security appliance, you can enable one or both of the following features:

- **Secure web proxy.** The appliance web proxy service monitors and scans web traffic for malicious content. When you enable the web proxy, you can configure it to be in transparent or explicit forward mode.
- **L4 Traffic Monitor.** The L4 Traffic Monitor detects and blocks rogue traffic across all ports and IP addresses. The L4 Traffic Monitor listens to network traffic that comes in over all ports and IP addresses on the appliance and matches domain names and IP addresses against entries in its own database tables to determine whether to allow outgoing traffic.

Tip — IronPort recommends enabling both the L4 Traffic Monitor and Web Proxy in the System Setup Wizard. If you need to disable both or one of these features, you can do so after initial setup from the web interface. If you do not enable one of the features in the System Setup Wizard and then need to enable it later, you must run the System Setup Wizard again, losing all configurations added to the appliance, such as group policies.

The features you enable determine how you deploy and physically connect the appliance to the network. For more information about how the features you enable affect appliance deployment, see “Preparing for Deployment” on page 14. For more information about the Ethernet ports used to physically connect the appliance to the network, see “Appliance Interfaces” on page 16.

Preparing for Deployment

Before installing the Web Security appliance, read through the following questions and use the responses to each question to help you decide how to deploy the appliance and where to locate it in your network. Each response includes a reference to a different section that covers the response in more detail.

1. Will you deploy the Web Security appliance as a transparent proxy or an explicit forward proxy?
   - **Explicit Forward Proxy.** Client applications, such as web browsers, are aware of the Web Proxy and must be configured to point to a single Web Security appliance. This deployment requires a connection to a standard network switch. When you deploy the Web Proxy in explicit forward mode, you can place it anywhere in the network. For more information, see “Deploying the Web Proxy in Explicit Forward Mode” on page 19.
   - **Transparent Proxy.** Clients applications are unaware of the Web Proxy and do not have to be configured to connect to the proxy. This deployment requires an L4 switch

or a WCCP v2 router. For more information, see "Deploying the Web Proxy in Transparent Mode" on page 20.

2. Does the network have an existing proxy?

   If yes, you must deploy the Web Security appliance downstream from an existing web proxy, meaning closer to the client. The Web Security appliance needs to sit between the internal clients and the existing proxy. This ensures that the appliance can effectively detect malware at the first hop and stop it before it gets to the existing proxy. The System Setup Wizard refers to this as an upstream proxy configuration.

   For more information, see "Using the Web Security Appliance in an Existing Proxy Environment" on page 26.

3. Will you enable the L4 Traffic Monitor?

   L4 Traffic Monitor deployment is independent of the Web Proxy deployment. You can connect the L4 Traffic Monitor to a network tap or the mirror/span port of a switch.

   For more information, see "Deploying the L4 Traffic Monitor" on page 27.

APPLIANCE INTERFACES

The Web Security appliance includes six physical Ethernet ports on the back of the system. Each Ethernet port corresponds to a different network interface. The Ethernet ports are grouped into the following types of network interfaces:

- **Management.** The Management interfaces include M1 and M2. However, only the M1 interface is enabled on the appliance. For more information, see "Management Interface" on page 16.
- **Data.** The Data interfaces include P1 and P2. Use the Data interfaces for Web Proxy data traffic. For more information, see "Data Interfaces" on page 16.
- **L4 Traffic Monitor.** The L4 Traffic Monitor interfaces include T1 and T2. Use these interfaces for monitoring and blocking L4 Traffic Monitor traffic. For more information, see "L4 Traffic Monitor Interfaces" on page 17.

Figure 2-1 shows the Ethernet ports on the back of the Web Security appliance blade.

Figure 2-1 Web Security Appliance Ethernet Ports

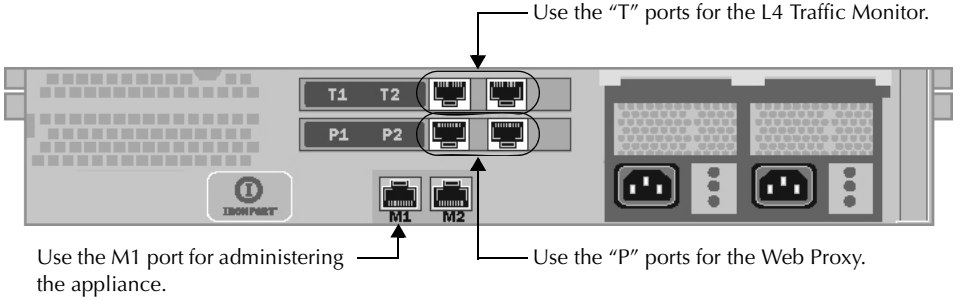

Management Interface

Use M1 to administer the appliance. Optionally, you can also configure the M1 interface to handle Web Proxy data traffic. You might want to use the M1 interface for data traffic if your organization does not use a separate management network. When M1 handles Web Proxy data traffic, neither of the data interfaces are enabled.

For more information about using the M1 port to set up and manage the appliance, see "Connecting a Laptop to the Appliance" on page 32.

For more information about configuring the network interfaces, see "Configuring Network Interfaces" on page 398.

Data Interfaces

The appliance uses the Data interfaces for Web Proxy data traffic. You can enable and use just the P1 port or both the P1 and P2 ports for data traffic.

- **P1 only enabled.** When only P1 is enabled, connect it to the network for both incoming and outgoing traffic.
- **P1 and P2 enabled.** When both P1 and P2 are enabled, you must connect P1 to the internal network and P2 toward the Internet.

Note — You can only enable and configure the P1 interface for data traffic in the System Setup Wizard. If you want to enable the P2 interface, you must use the `ifconfig` command after finishing the System Setup Wizard. For more information about configuring the P2 interface, see "Configuring Network Interfaces" on page 398.

How you physically connect the data interfaces to the network depends on how you deploy the appliance. For more information, see "Deploying the Web Proxy in Explicit Forward Mode" on page 19 and "Deploying the Web Proxy in Transparent Mode" on page 20.

L4 Traffic Monitor Interfaces

The appliance uses the T1 and T2 interfaces for listening to traffic on all TCP ports. You can connect just T1 or both T1 and T2 using an Ethernet cable, depending on whether you use simplex or duplex communication.

- **T1 only connected (duplex).** When you configure the appliance to use duplex communication, connect T1 to the network so it receives all incoming and outgoing traffic.
- **T1 and T2 connected (simplex).** When you configure the appliance to use simplex communication, connect T1 to the network so it receives all outgoing traffic (from the clients to the Internet), and connect T2 to the network so it receives all incoming traffic (from the Internet to the clients).

For more information about how to connect the L4 Traffic Monitor ports to the network, see "Deploying the L4 Traffic Monitor" on page 27.

Example Deployment

Figure 2-2 on page 18 shows a sample deployment scenario with both the Web Proxy and L4 Traffic Monitor enabled. In this example, the Web Proxy is deployed in transparent mode and only the P1 port is connected to either a L4 switch or a WCCP router.

Figure 2-2 Web Security Appliance Deployment Scenario

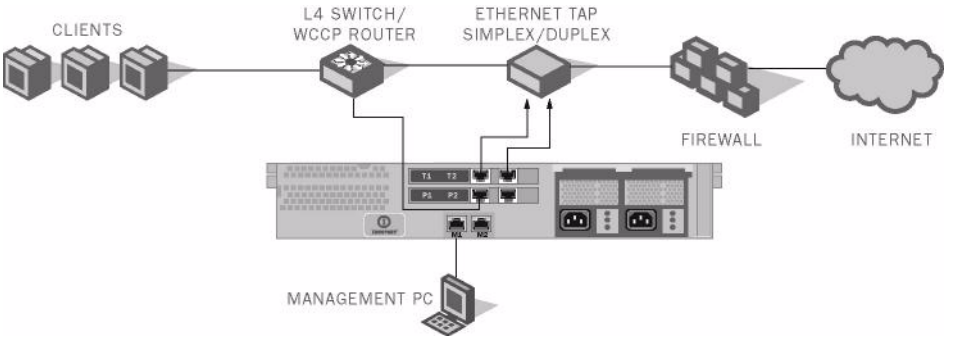

DEPLOYING THE WEB PROXY IN EXPLICIT FORWARD MODE

When the appliance is configured as an explicit forward proxy, client applications must be configured to direct its traffic to the appliance. When you want to configure the Web Proxy in explicit forward mode, you must configure the following components:

- Client applications
- Appliance ports

Tip — If your organization needs to use explicit forward mode now, but might need transparent mode in the future, IronPort recommends enabling Transparent mode in the System Setup Wizard and then choosing L4 switch as the connection type. If you do not have an L4 switch, you can connect the appliance to the network normally and the appliance will work in explicit forward mode. To use transparent mode in the future, you can connect the appliance to an L4 switch and it will work in transparent mode without needing to run the System Setup Wizard again.

Configuring Client Applications

You must configure the all client applications, such as web browsers, used on the network to point to the Web Proxy. You can configure each client in the following ways:

- **Manual.** Configure each client application to point the appliance Web Proxy by specifying the appliance host name or IP address and the port number, such as 3128, used for listening to data traffic.
- **Automatic.** Configure each client application to use a PAC file to detect the appliance Web Proxy automatically. Then you can edit the PAC file to specify the appliance Web Proxy information. For more information, see “Working with PAC Files” on page 65.

Connecting Appliance Interfaces

You can connect the P1 interface or both the P1 and P2 interfaces to a network switch using an Ethernet cable. You do not need special hardware, such as a particular switch or router. For more information about how to connect the data interfaces (P1 and P2), see “Data Interfaces” on page 16.

Testing an Explicit Forward Configuration

If you want to test an explicit forward proxy configuration, you can separate and forward traffic from a subset of your network infrastructure. To individually test this configuration, clients can forward traffic to the appliance from one web browser and connect to the internet using another web browser. This method also ensures an alternate path to the Internet while testing.

DEPLOYING THE WEB PROXY IN TRANSPARENT MODE

When the appliance is configured as a transparent proxy, client applications are not aware that their traffic gets redirected to the appliance, and they do not need to be configured to point to the appliance. To deploy the appliance in this mode, you need one of the following types of hardware to transparently redirect web traffic to the appliance:

- **WCCP v2 router.** When you specify a WCCP router, you need to configure additional settings on the appliance. For more information about using the appliance with a WCCP router, see "Connecting the Appliance to a WCCP Router" on page 21.
- **Layer 4 switch.** When you specify an L4 switch, you only need to specify that the appliance is connected to an L4 switch when you configure the appliance. You do not need to configure anything else on the appliance.

Typically, you configure the appliance to use an L4 switch or a WCCP v2 router during initial system setup. However, you can configure it to use either an L4 switch or a WCCP v2 router anytime after initial setup on the Network > Transparent Redirection page. For more information about the Network > Transparent Redirection page, see "Configuring Transparent Redirection" on page 402.

Connecting Appliance Interfaces

When you configure the Web Proxy in transparent mode, you can connect the P1 port or both the P1 and P2 ports to an L4 switch or WCCP router using an Ethernet cable. For more information about how to connect the data interfaces (P1 and P2), see "Data Interfaces" on page 16.

CONNECTING THE APPLIANCE TO A WCCP ROUTER

When you connect the appliance to a WCCP router, you must perform the following tasks:

1. You must create at least one WCCP service on the appliance. For more information, see "Configuring the Web Security Appliance" on page 21.
2. After you create a WCCP service, you must configure the router to work with the Web Security appliance. For more information, see "Configuring the WCCP Router" on page 21.

You can also connect an appliance to multiple WCCP routers. For more information, see "Working with Multiple Appliances and Routers" on page 25.

Configuring the Web Security Appliance

A WCCP service is an appliance configuration that defines a service group to a WCCP v2 router. It includes information such as the service ID and ports used. Service groups allow a web proxy to establish connectivity with a WCCP router and to handle redirected traffic from the router.

Create WCCP services on the Network > Transparent Redirection page. The WCCP services you create determine how you configure the WCCP routers. For more information about creating WCCP services, see "Adding and Editing a WCCP Service" on page 405.

Note — You can only create WCCP services in the web interface after initial setup. You cannot create WCCP services in the System Setup Wizard.

Configuring the WCCP Router

After you create at least one WCCP service in the Web Security appliance, you can configure the WCCP router(s) in the network.

Use the following syntax for enabling WCCP on the router:

```
ip wccp version 2
ip wccp service_group
interface interface_type_number
ip wccp service_group redirect direction
ip wccp service_group password password
```

Enter one of the following values for the *service_group* variable:

- **web-cache.** Enter "web-cache" when the appliance WCCP service uses the standard service.
- **Service ID number.** Enter a number from 0 to 255 when the appliance WCCP service uses a dynamic service ID. The number should match the service ID number used in the appliance.

Table 2-1 describes each part of the WCCP configuration syntax for enabling WCCP on the router.

Table 2-1 WCCP Router Configuration Syntax for Enabling the Router

WCCP Configuration	Description
`ip wccp version 2`	Defines the version of WCCP to use on the router. You must specify version 2 to work with the Web Security appliance. This command is required.
`ip wccp` *`service_group`* `password` *`password`*	Specifies a service group to enable on the router. It also enables the WCCP service on the router. This command is required.
`interface` *`interface_type_number`*	Specifies an interface to configure and enters interface configuration mode. Enter the interface number for the *interface_type_number* variable. This command is required.
`ip wccp` *`service_group`* `redirect` *`direction`*	Enables WCCP redirection on the specified interface. Enter one of the following values for the *direction* variable: • **in.** Use `in` when you want the router to redirect packets as they enter the router. • **out.** Use `out` when you want the router to redirect packets right before they leave the router. This command is required.
`ip wccp` *`service_group`* `password` *`password`*	Sets a password on the router for the specified service group. This command is only required when the WCCP service defined on the appliance has password security enabled.

You can also configure a WCCP router to perform other tasks, such as the following:

- Configure the router from exclude redirecting traffic received on a particular interface.
- If the network uses multiple Web Security appliances, you can configure the router to determine which traffic should be directed to which appliance by using an access list. You might want to redirect only some of the network traffic to the appliance if you are evaluating the Web Security appliance.

For more detailed information about configuring WCCP routers, go to the following location:

`http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a008030c778.html`

Note — The Web Security appliance does not support using a multicast address in the WCCP service group. To use multiple routers in a service group, you must specify the IP address of

each router in the service group and configure each router separately. You cannot register a router to a multicast address.

Example WCCP Configurations

This section shows some sample WCCP services defined in the appliance and the corresponding WCCP configuration you should use to configure the router that connects to the appliance.

Example 1

Suppose you have the WCCP service shown in Figure 2-3.

Figure 2-3 Example WCCP Service — Standard Service, No Password Required

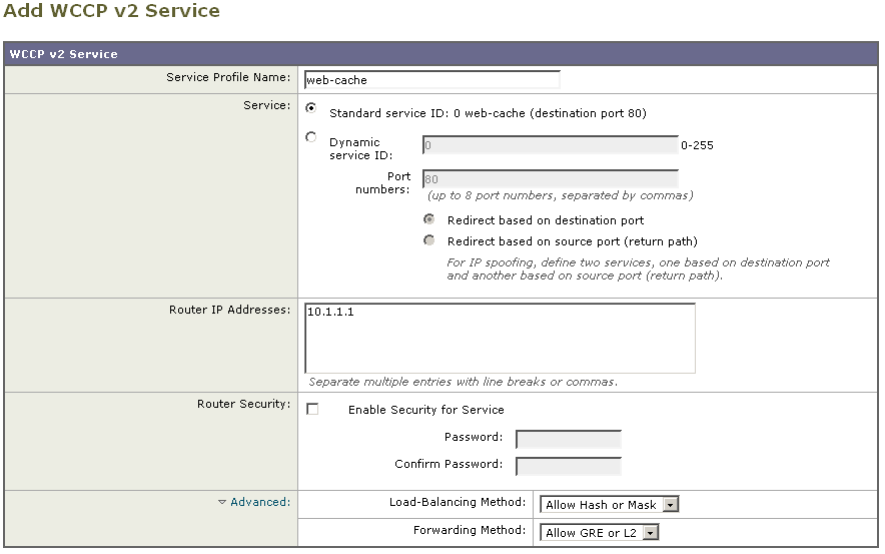

In this example, the WCCP service defines the standard service group (also known as a well known service group). The redirection basis is on the destination port by default. Also suppose in this example that you want to configure the ethernet1 interface on the router for this service group.

Use the following WCCP configuration for this example:

```
ip wccp version 2
ip wccp web-cache
interface ethernet1
ip wccp web-cache redirect in
```

Example 2

Figure 2-4 shows a dynamic service you might create when IP spoofing is enabled and the WCCP service shown in Figure 2-3 on page 23 is defined.

Figure 2-4 Example WCCP Service — Dynamic Service for IP Spoofing

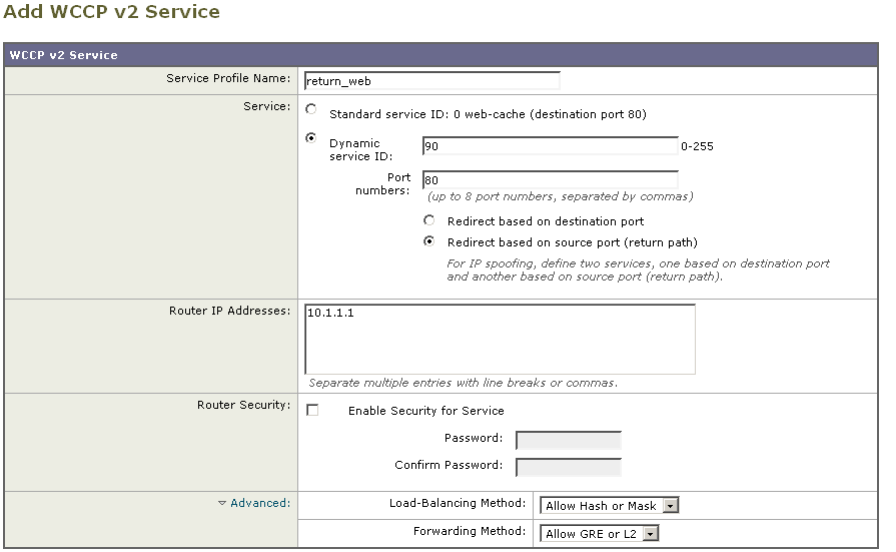

In this example, the WCCP service defines a dynamic service group with service ID of 90. The redirection basis is on the source port so it can be used for the return path with IP spoofing enabled. Suppose in this example that you want to configure the ethernet0 interface on the router for this service group.

Use the following WCCP configuration for this example:

```
ip wccp version 2
ip wccp 90
interface ethernet0
ip wccp 90 redirect out
```

For more information about enabling IP spoofing when using a WCCP router, see “IP Spoofing when Using WCCP” on page 404.

Example 3

Suppose you have the WCCP service shown in Figure 2-5.

Figure 2-5 Example WCCP Service — Dynamic Service, Password Required

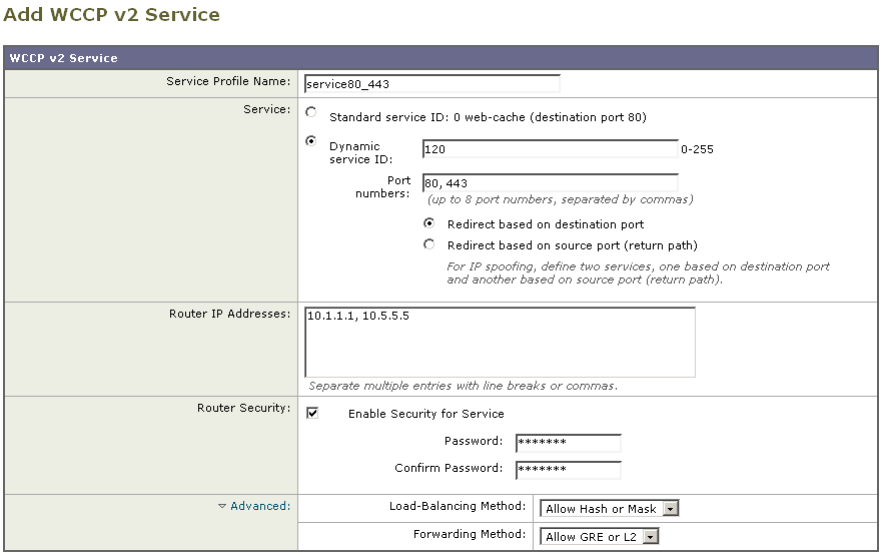

In this example, the WCCP service defines a dynamic service group with service ID of 120. The redirection basis is on the destination port, and it has enabled a password for this service group of "admin99" (hidden in the appliance configuration). Also suppose in this example that you want to configure the ethernet0 interface on the router for this service group.

Use the following WCCP configuration for this example:

```
ip wccp version 2
ip wccp 120
interface ethernet0
ip wccp 120 redirect in
ip wccp 120 password admin99
```

Working with Multiple Appliances and Routers

When you connect one or more Web Security appliances to one or more WCCP routers, you have a cluster. You can include up to 32 appliances and up to 32 routers in a cluster. You must configure all appliances and routers in a cluster to communicate with each other.

USING THE WEB SECURITY APPLIANCE IN AN EXISTING PROXY ENVIRONMENT

The Web Security appliance is a proxy-compatible device, and is easily deployed within an existing proxy environment. However, you must place the appliance downstream from existing proxy servers, meaning closer to the clients. Positioning the appliance between internal clients and an existing proxy server enables the appliance to detect and stop malware before it reaches the existing proxy server.

You can configure the appliance to work with an existing, upstream proxy in the System Setup Wizard or after the initial setup in the web interface. Use the Network > Upstream Proxies page to enable an upstream proxy or to modify existing settings.

When configuring an upstream proxy, you specify whether the existing proxy is in transparent or explicit forward mode.

Transparent Upstream Proxy

If a transparent upstream proxy uses client IP addresses to manage user authentication and access control, you must enable IP spoofing on the Web Security appliance to send client IP addresses to the upstream proxy. Use the Security Services > Proxy Settings page to enable IP spoofing.

When you enable IP spoofing and connect the appliance to a WCCP router, you must create at least two WCCP services. For more information about configuring WCCP services when you enable IP spoofing, see “IP Spoofing when Using WCCP” on page 404.

Explicit Forward Upstream Proxy

If the upstream proxy is in explicit forward mode, consider the following rules and guidelines:

- You must enter the IP address or host name and port of the upstream proxy.
- Consider whether the host name of the upstream proxy resolves to multiple IP addresses. The Web Security appliance only queries the DNS server for the IP address at startup. If an IP address is added or removed from that host name, the proxy must restart to resolve and add the host name to the new set of IP addresses.
- If the upstream proxy manages user authentication or access control using proxy authentication, you must enable the X-Forwarded-For header to send the client host header to the upstream proxy. Use the Security Services > Proxy Settings page to enable the X-Forwarded-For header setting.
- If the upstream proxy manages client traffic using a PAC file or a login script, you must update these files to use the IP address or host name of the Web Security appliance.

DEPLOYING THE L4 TRAFFIC MONITOR

L4 Traffic Monitor (L4TM) deployment is independent of the Web Proxy deployment. When connecting and deploying the L4 Traffic Monitor, consider the following:

- **Physical connection.** You can choose how to connect the L4 Traffic Monitor to the network. For more information, see "Connecting the L4 Traffic Monitor" on page 27.
- **Network address translation (NAT).** When configuring the L4 Traffic Monitor, connect it at a point in your network where it can see as much network traffic as possible before getting out of your egress firewall and onto the Internet. It is important that the L4 Traffic Monitor be 'logically' connected after the proxy ports and before any device that performs network address translation (NAT) on client IP addresses.
- **L4 Traffic Monitor action setting.** The default setting for the L4 Traffic Monitor is monitor only. After setup, if you configure the L4 Traffic Monitor to monitor and block suspicious traffic, ensure that the L4 Traffic Monitor and the Web Proxy are configured on the same network so that all clients are accessible on routes that are configured for data traffic.

Connecting the L4 Traffic Monitor

You can connect the L4 Traffic Monitor to the network in any of the following ways:

- **Network tap.** When you use a network tap, you can choose the following communication types:
  - **Simplex.** This communication type uses one cable for all traffic between clients and the appliance, and one cable for all traffic between the appliance and external connections. Connect port T1 to the network tap so it receives all outgoing traffic (from the clients to the Internet), and connect port T2 to the network tap so it receives all incoming traffic (from the Internet to the clients).
  - **Duplex.** This mode uses one cable for all incoming and outgoing traffic. You can use half- or full-duplex Ethernet connections. Connect port T1 to the network tap so it receives all incoming and outgoing traffic.

  Note — IronPort recommends using simplex when possible because it can increase performance and security.

- **Span/mirror port of an L2 switch.** Connecting is similar to a simplex or duplex tap, depending on whether the connection uses two separate devices or one device.
- **Hub.** Choose duplex when you connect the L4 Traffic Monitor to a hub.

Regardless of how the appliance is connected to the network, you must configure the wiring type. For more information, see "Configuring an L4 Traffic Monitor Wiring Type" on page 28.

For more information about the T1 and T2 ports, see "Appliance Interfaces" on page 16.

Note — Use a network tap instead of the span/mirror port of a switch when possible. Network taps use hardware to move packets to the L4 Traffic Monitor and span and mirror ports of a

switch use software to move packets. Hardware solutions move packets with better performance than software solutions and are less likely to drop packets in the process.

Configuring an L4 Traffic Monitor Wiring Type

Typically, the L4 Traffic Monitor wiring type is configured during system setup. However, you can configure the wiring type after running the System Setup Wizard on the Network > Interfaces page. Click **Edit Settings** and select a wiring type for the T1 and T2 ports.

Figure 2-6 L4 Traffic Monitor Wiring Types

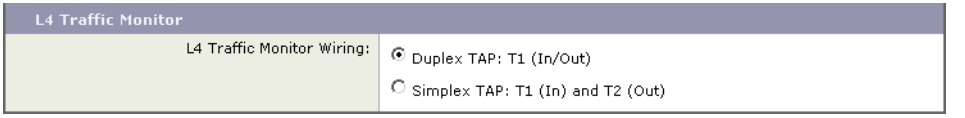

PHYSICAL DIMENSIONS

The following physical dimensions apply to the **IronPort S660 and S360** Web Security appliances:

- Height: 8.656 cm (3.40 inches)
- Width: 48.26 cm (19.0 inches) with rails installed (without rails, 17.5 inches)
- Depth: 75.68cm (29.79 inches)
- Weight: maximum 25.6 kg (56.6 pounds)

The following physical dimensions apply to the **IronPort S160** Web Security appliance:

- Height: 4.2 cm (1.68 inches)
- Width: 48.26 cm (19.0 inches) with rails installed (without rails, 17.5 inches)
- Depth: 57.6 cm (22.7 inches)
- Weight: maximum 9.8 kg (21.6 pounds)

CHAPTER 3

Installation and Configuration

This chapter contains the following topics:

BEFORE YOU BEGIN

To use the Web Security appliance, you must run the System Setup Wizard. However, first you must do some steps to prepare the appliance for the System Setup Wizard.

For more information about preparing the appliance for installation, see the Web Security appliance *QuickStart Guide*. You can find this guide and other useful information about the IronPort Web Security appliance Support Portal:

`https://supportportal.ironport.com/irppcnctr/srvcd?u=http://secure-support.soma.ironport.com`

Complete the following tasks before you run the System Setup Wizard:

- **Deployment.** Decide how you are going to configure the appliance within your network. For details, see "Deployment" on page 13.
- **Laptop network connection.** Configure your laptop's network connection to use an IP address on the same subnet as the Web Security appliance (192.168.42.xx). For details, see "Connecting a Laptop to the Appliance" on page 32.
- **Appliance physical connections.** Plug the Ethernet cables into the appropriate ports on the back panel of the appliance. For details, see "Connecting the Appliance to the Network" on page 32.
- **Setup information.** Once you know how you will install the appliance in your network, gather all the information, such as IP addresses, necessary for the System Setup Wizard. For details, see "Gathering Setup Information" on page 33.
- **Existing proxy server.** If you plan to use the Web Security appliance in a network that has an existing proxy server, you must locate it downstream from other proxy servers. Also, after you finish the initial setup of the appliance, you must configure it to work with the existing proxy server. For more information about deploying the appliance in a network with an existing proxy, see "Using the Web Security Appliance in an Existing Proxy Environment" on page 26.

Connecting a Laptop to the Appliance

In order to run the System Setup Wizard the first time, you must connect a computer, such as a laptop, to the appliance. To connect to the appliance, the laptop subnet must be the same as the appliance subnet. The Management ports are labeled M1 and M2. The Web Security appliance only uses the M1 Management port. It does not use M2.

Configure the laptop IP address so it is on the same subnet as the appliance (192.168.42.xx). Then, connect the laptop to the M1 port on the back of the appliance.

Connecting the Appliance to the Network

You must plug the Ethernet cables into the appropriate ports on the back panel of the appliance. For more information about the Ethernet ports on the appliance, see "Appliance Interfaces" on page 16.

How you deploy the appliance determines which Ethernet cables to plug in where:

- **Web proxy in transparent mode.** If you want to use one proxy port for all traffic, connect port P1 to an L4 switch or a WCCP router using an Ethernet cable. If you want to use two proxy ports for traffic, connect port P2 to an L4 switch or a WCCP router using an Ethernet cable, and connect port P1 to the internal network.

  For more information about deploying the Web Proxy in transparent mode, see "Deploying the Web Proxy in Transparent Mode" on page 20.

  Note — When you configure the proxy in transparent mode and connect it to a WCCP router, you must configure the appliance after you run the System Setup Wizard to create at least one WCCP service. For more information about creating WCCP services, see "Adding and Editing a WCCP Service" on page 405.

- **Web proxy in explicit forward mode.** If you want to use one proxy port for all traffic, connect port P1 to a network switch using an Ethernet cable. If you want to use two proxy ports for traffic, connect port P2 to a network switch using an Ethernet cable, and connect port P1 to the internal network.

  For more information about deploying the Web Proxy in explicit forward mode, see "Deploying the Web Proxy in Explicit Forward Mode" on page 19.

- **L4 Traffic Monitor.** Connect the Traffic Monitor ports to the Ethernet tap according to the tap communication type:

  - **Ethernet tap using simplex.** Connect port T1 to the Ethernet tap so it receives all outgoing traffic (from the clients to the Internet), and connect port T2 to the Ethernet tap so it receives all incoming traffic (from the Internet to the clients).
  - **Ethernet tap using duplex.** Connect port T1 to the Ethernet tap so it receives all incoming and outgoing traffic.

  For more information about deploying the L4 Traffic Monitor, see "Deploying the L4 Traffic Monitor" on page 27.

Gathering Setup Information

Once you know how you will install the appliance in your network, you can gather the necessary information, such as IP addresses, to enter in the System Setup Wizard. You can use the worksheet in Table 3-1 to write down the configuration options you decide on. Then, when you run the System Setup Wizard, you can use the information you enter in the worksheet to configure the initial setup.

Table 3-1 System Setup Worksheet

Deployment Options	
Web Proxy:	Enable / Disable

Table 3-1 System Setup Worksheet (Continued)

If Enabled Proxy:	Transparent with L4 switch / Transparent with WCCP router / Explicit forward proxy
Layer 4 Traffic Monitor:	Enable / Disable
If Enabled Layer 4 Traffic Monitor:	Simplex network tap / Duplex network tap
Network Context	
Is there another proxy on the network:	No / Yes in Transparent Mode / Yes in Forward Mode
Other Proxy in Forward Mode IP Address:	
Other Proxy in Forward Mode Port:	
Network Settings	
Default System Host name:	See "DNS Support" on page 36 for more information.
DNS Servers:	Internet root DNS servers / organization DNS servers
Organization DNS Servers: (maximum 3)	1. 2. 3.
Network Time Protocol Server:	
Time Zone Region:	
Time Zone Country:	
Time Zone / GMT Offset:	
Interface Settings	
Management Port	
IP Address:	
Network Mask:	
Host Name:	
Data Port	

Table 3-1 System Setup Worksheet (Continued)

IP Address:	
Network Mask:	
Host Name:	
Note: The Web Proxy can share the Management interface. If configured separately, the Data interface IP address and the Management interface IP address cannot share the same subnet.	
Routes	
Management Traffic	
Default Gateway:	
Static Route Table Name:	
Static Route Table Destination Network:	
Static Route Table Gateway:	
Data Traffic	
Default Gateway:	
Static Route Table Name:	
Static Route Table Destination Network:	
Static Route Table Gateway:	
Transparent Routing Device	
Device Type:	Layer 4 switch / WCCP Router
Note: When you connect the appliance to a WCCP router, you must configure the Web Security appliance to create WCCP services after you run the System Setup Wizard. For more information about creating WCCP services, see “Adding and Editing a WCCP Service” on page 405.	
Administrative Settings	
Administrator Password:	
Email System Alerts To:	
SMTP Relay Host:	(optional)

Table 3-1 System Setup Worksheet (Continued)

AutoSupport:	Enable / Disable
Security Services	
IP Spoofing:	Enable / Disable
L4 Traffic Monitor:	Monitor only / Block
IronPort URL Filtering:	Enable / Disable
Web Reputation Filters:	Enable / Disable
Malware and Spyware Scanning:	Enable Webroot / Enable McAfee / Enable both
Action for Detected Malware:	Monitor only / Block
Action for Unscannable Transactions:	Monitor only / Block
SenderBase Network Participation:	Enable / Disable
Participation Level:	Limited / Standard

DNS Support

To connect to the management interface using a host name (http://hostname:8080), you must add the appliance host name and IP address to your DNS server database.

SYSTEM SETUP WIZARD

The IronPort AsyncOS for Web operating system provides a browser-based wizard to guide you through initial system configuration. This System Setup Wizard prompts you for basic initial configuration, such as network configuration and security settings. The System Setup Wizard is located on the System Administration tab.

You must run the System Setup Wizard when you first install the Web Security appliance. After you finish the System Setup Wizard, the appliance is ready to monitor web traffic. However, you may want to make more custom configurations to the appliance that the System Setup Wizard does not cover. For more information about configuration options, see most of the other chapters in this guide.

Before you run the System Setup Wizard, see "Before You Begin" on page 32 to verify you have all the information you need to configure the appliance. Having this information prepared ahead of time can reduce the amount of time required to complete the initial setup. You should also read the *QuickStart Guide* for more information about product setup.

WARNING: Running the System Setup Wizard completely reconfigures the Web Security appliance and resets the administrator password. Only use the System Setup Wizard the first time you install the appliance, or if you want to completely overwrite the existing configuration. Running the System Setup Wizard after the appliance is already configured can also interrupt client access to the web. If you choose to run the System Setup Wizard after performing an initial setup, use the System Administration > Configuration File pages to print a configuration summary and archive the current configuration file.

WARNING: The IronPort appliance ships with a default IP address of 192.168.42.42 on the Management interface (port). Before connecting the appliance to the network, ensure that no other device on the network has the same IP address.

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

The System Setup Wizard includes the following tabs where you enter configuration information:

- **Start.** For details, see "Step 1. Start" on page 38.
- **Deployment.** For details, see "Step 2. Deployment" on page 38.
- **Network.** For details, see "Step 3. Network" on page 43.
- **Security.** For details, see "Step 4. Security" on page 50.
- **Review.** For details, see "Step 5. Review" on page 52.

Accessing the System Setup Wizard

To access the System Setup Wizard, open a browser and enter the IP address of the Web Security appliance. The first time you run the System Setup Wizard, use the default IP address:

```
http://192.168.42.42
```

The appliance login screen appears. Enter the username and password to access the appliance. By default, the appliance ships with the following username and password:

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

Note — Your session will time out if it is idle for over 30 minutes or if you close your browser without logging out. If this happens, you must re-enter the username and password.

Step 1. Start

When you first start the System Setup Wizard, it displays an end user license agreement.

1. Accept the terms of the agreement by clicking the check box at the bottom of the page.

Figure 3-1 System Setup Wizard — Start Tab

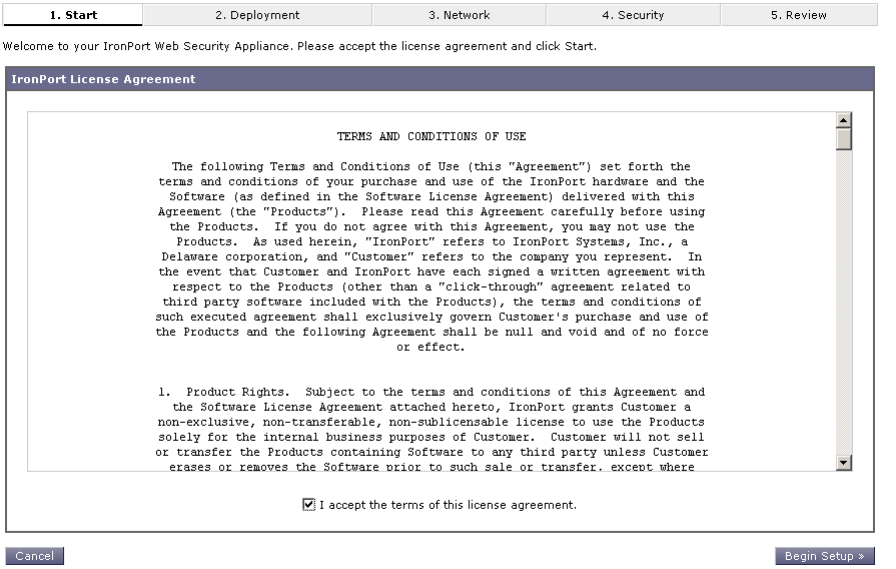

2. Click **Begin Setup** to continue.

   The Deployment tab appears.

Step 2. Deployment

The System Setup Wizard continues to the Deployment tab. The Deployment tab contains multiple pages that prompt you to enter information.

1. Verify that the Web Security Appliance Functions page appears.

Figure 3-2 shows the Web Security Appliance Functions page of the Deployment tab.

Figure 3-2 System Setup Wizard — Deployment Tab, Web Security Appliance Functions Page

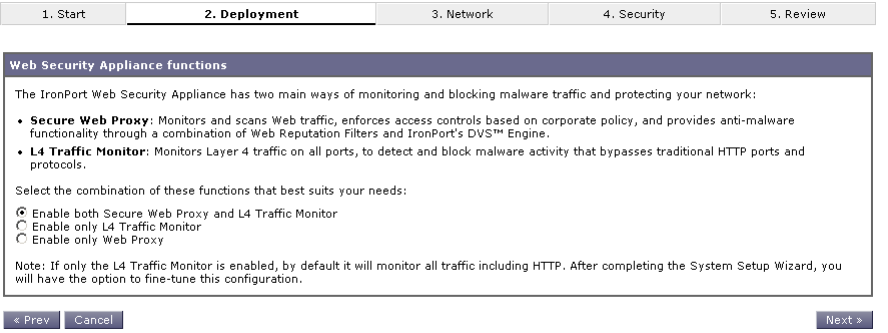

2. Choose the Web Security Appliance Functions options.

   Table 3-2 describes the Web Security Appliance Function options.

Table 3-2 Web Security Appliance Functions Options in System Setup Wizard

Option	Description
Enable both Secure Web Proxy and L4 Traffic Monitor	Configures the Web Security appliance to enable both the Web Proxy and L4 Traffic Monitor features. For more information about choosing which features to enable, see “Deployment Overview” on page 14.
Enable only L4 Traffic Monitor	Configures the Web Security appliance to enable only the L4 Traffic Monitor. For more information about choosing which features to enable, see “Deployment Overview” on page 14. The L4 Traffic Monitor detects rogue traffic across all network ports and stops malware attempts to bypass port 80. For more information about the L4 Traffic Monitor, see “L4 Traffic Monitor” on page 295.

Table 3-2 Web Security Appliance Functions Options in System Setup Wizard (Continued)

Option	Description
Enable only Web Proxy	Configures the Web Security appliance to enable only the Web Proxy service. For more information about choosing which features to enable, see "Deployment Overview" on page 14. The Web Proxy service monitors and controls traffic that originates from clients on the internal network. Typically, the Web Proxy is deployed between clients and the firewall where it intercepts requests for content from clients to servers. For more information about the Web Proxy, see "Web Proxy Services" on page 55.

3. Click **Next**.

   The Network Context page appears.

   Figure 3-3 System Setup Wizard — Deployment Tab, Network Context Page

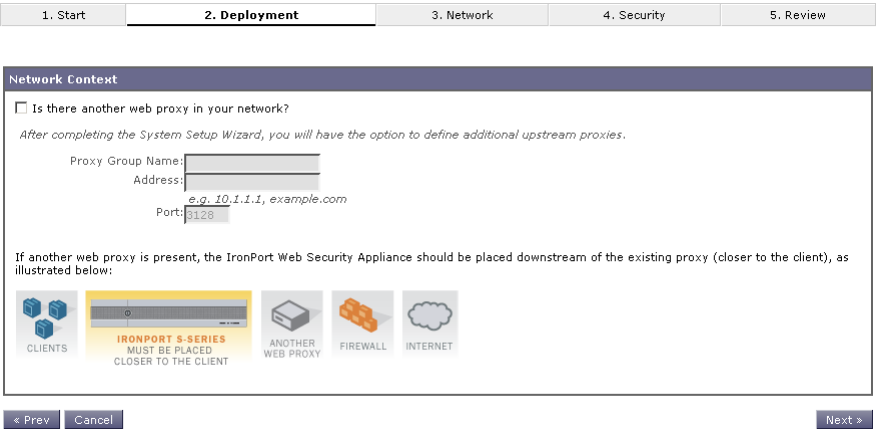

4. Configure the Network Context options by indicating whether or not there exists another proxy server on the network.

   Note — You can configure the Web Security appliance to interact with multiple proxy servers on the network after you run the System Setup Wizard. For more information about configuring external proxy servers, see "Working with External Proxies Overview" on page 138.

5. If there is an external proxy server on the network, configure the proxy settings.

Table 3-3 describes the proxy settings.

Table 3-3 Network Context Options in System Setup Wizard

Option	Description
Proxy group name	Choose a name for the proxy group.
Address	Enter the address of the proxy server in your organization network.
Port	The port number of the proxy server in your organization network.

The System Setup Wizard creates a proxy group with the information you provide in Table 3-3. You can edit the proxy group later to include additional proxy servers and to configure load balancing options. You can also create additional proxy groups after system setup.

Note — When you use the Web Security appliance in a network that contains another proxy server, you must place the Web Security appliance downstream from the proxy server, closer to the clients.

6. Click **Next**.

   If you configured the Web Security appliance as a web proxy on the Web Security Appliance Functions page, then the Proxy Mode page appears. If you did you not configure the appliance as web proxy, then skip to step 9 on page 42.

   Figure 3-4 System Setup Wizard — Deployment Tab, Proxy Mode Page

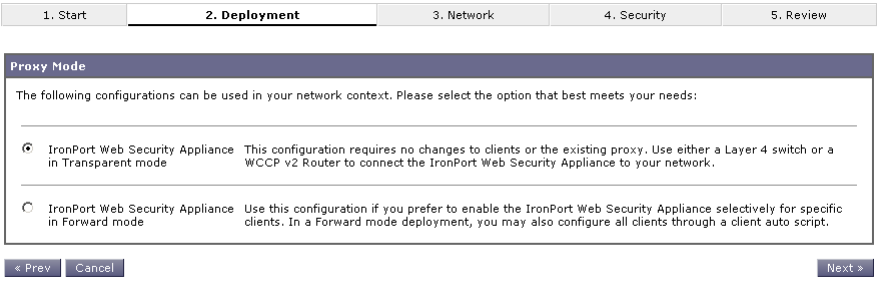

7. Configure the Proxy Mode options.

Table 3-4 describes the Proxy Mode options.

Table 3-4 Proxy Mode Options in System Setup Wizard

Option	Description
IronPort Web Security Appliance in Transparent mode	Configures the Web Security appliance as a transparent proxy server. Choose this option when you do not want to configure client applications, such as web browsers, to be aware of the proxy service in the appliance. For more information about choosing which mode to use, see “Deployment Overview” on page 14. For more information about configuring the appliance as a proxy server, see “Web Proxy Services” on page 55.
IronPort Web Security Appliance in Forward mode	Configures the Web Security appliance as a forward proxy server. Choose this option when you want the appliance to act on behalf of client web browsers to handle requests for web servers. When you choose this option, users must configure their web browsers to point to a single Web Security appliance. For more information about choosing which mode to use, see “Deployment Overview” on page 14. For more information about configuring the appliance as a proxy server, see “Web Proxy Services” on page 55.

8. Click **Next**.

   The Deployment Summary page appears.

   Figure 3-5 System Setup Wizard — Deployment Tab, Summary Page

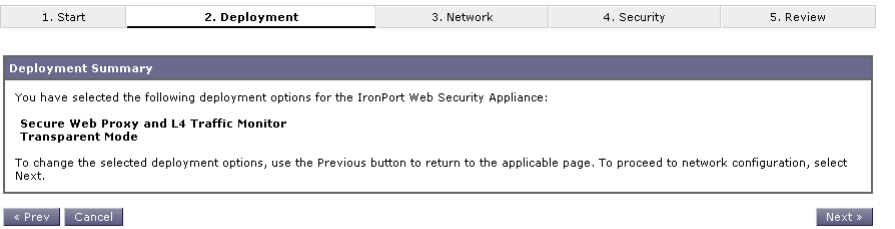

9. Verify that the Deployment options are correct, and click **Next**.

   The Network tab appears.

Step 3. Network

On the Network tab, you configure appliance system properties, such as the appliance host name and time zone. The first page of the Network tab is the System Configuration page.

1. Verify that you are viewing the System Configuration page.

Figure 3-6 System Setup Wizard — Network Tab, System Configuration

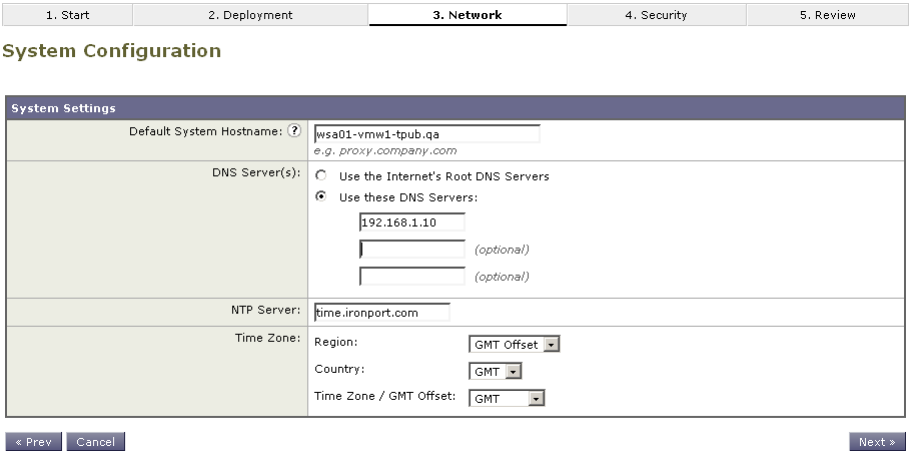

2. Configure the System Configuration options.

Table 3-5 describes the System Configuration options.

Table 3-5 System Configuration Options in System Setup Wizard

Option	Description
Default System Hostname	The fully-qualified hostname for the Web Security appliance. This name should be assigned by your network administrator. This hostname is used to identify the appliance in system alerts.

Table 3-5 System Configuration Options in System Setup Wizard (Continued)

Option	Description
DNS Server(s): Use the Internet's Root DNS Servers	Configures the appliance to use the Internet root DNS servers for domain name service lookups. You might choose this option when the appliance does not have access to DNS servers on your network. The appliance requires access to a working DNS server in order to perform DNS lookups for incoming connections. If you cannot specify a working DNS server that the appliance can reach while you set up the appliance, you can configure it to use the Internet root DNS servers or temporarily assign the IP address of the Management interface so that you can complete the System Setup Wizard. For more information about configuring DNS settings, see "Configuring DNS Server(s)" on page 393.
DNS Server(s): Use these DNS Servers	Specifies local DNS servers for domain name service lookups. You must enter at least one DNS server, and up to three total. You can choose to use the Internet root DNS servers or specify your own DNS servers. For more information about configuring DNS settings, see "Configuring DNS Server(s)" on page 393.
NTP Server	Uses a Network Time Protocol (NTP) server to synchronize the system clock with other servers on the network or the Internet. By default, the IronPort Systems time server (time.ironport.com) is entered.
Time Zone	Sets the time zone on the IronPort 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 using the GMT offset. For more information about the GMT offset, see "Selecting a Time Zone" on page 409.

3. Click **Next**.

   The Network Interfaces and Wiring page appears.

   The Web Security appliance has network interfaces that are associated with the physical ports on the machine.

Figure 3-7 System Setup Wizard — Network Tab, Network Interfaces and Wiring Page

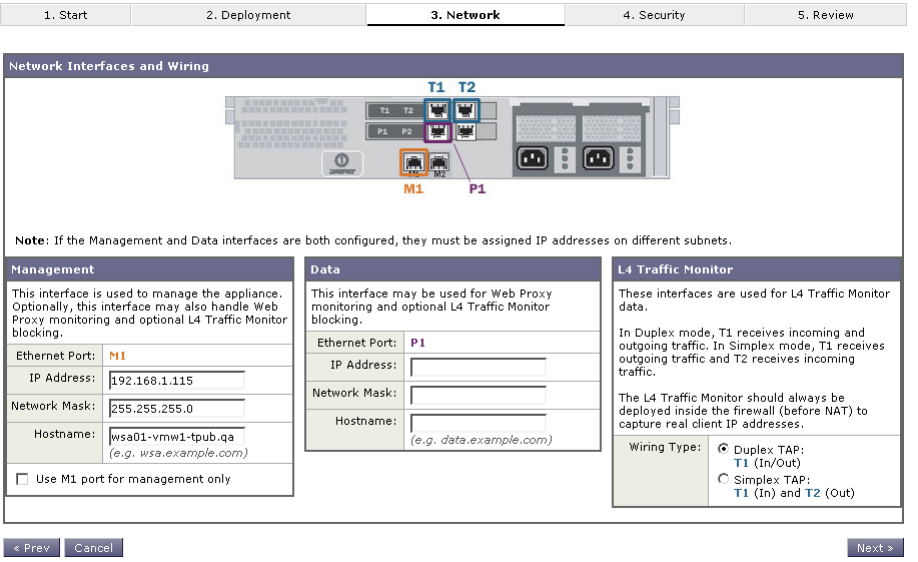

4. Configure the Network Interfaces and Wiring options.

   The appliance has network interfaces that are associated with the physical ports on the machine. Table 3-6 describes the Network Interfaces and Wiring options.

Table 3-6 Network Interfaces and Wiring Options in System Setup Wizard

Option	Description
Management	Enter the IP address, network mask, and hostname to use to manage the Web Security appliance. Enter an IP address that exists on your management network. By default, the appliance uses the M1 interface for both management and proxy (data) traffic (the "Use M1 port for management only" check box is *disabled*). However, optionally, you can use the M1 interface for only management traffic by enabling the "Use M1 port for management only" check box. You might want to do this if your organization uses a separate management network. This can increase security by ensuring no proxy traffic can reach the appliance on management interface. When you use M1 for management traffic only, you must configure at least one data interface for proxy traffic. Also, you must define different routes for management and data traffic.

Table 3-6 Network Interfaces and Wiring Options in System Setup Wizard (Continued)

Option	Description
Data	Enter the IP address, network mask, and hostname to use for data traffic. If you configure the M1 interface for management traffic only, you must configure the P1 interface for data traffic. However, you can configure the P1 interface even when the M1 interface is used for both management and data traffic. You can use the Data interface for Web Proxy monitoring and optional L4 traffic monitoring. You can also configure this interface to support outbound services, such as DNS, software upgrades, NTP, and traceroute data traffic. **Note:** You can only enable and configure the P1 network interface for data traffic in the System Setup Wizard. If you want to enable the P2 interface, you must use the `ifconfig` command after finishing the System Setup Wizard. For more information about configuring the P2 interface, see "Configuring Network Interfaces" on page 398.
L4 Traffic Monitor	Choose the type of wired connections plugged into the "T" interfaces: • **Duplex TAP.** Choose Duplex TAP when the T1 port receives both incoming and outgoing traffic. You can use half- or full-duplex Ethernet connections. • **Simplex TAP.** Choose Simplex TAP when you connect the T1 port to the internal network (traffic flows from the clients to the Internet) and you connect the T2 port to the external network (traffic flows from the Internet to the clients). IronPort recommends using Simplex when possible because it can increase performance and security.

5. Click **Next**.

   The Routes for Management and Data Traffic page appears.

Figure 3-8 System Setup Wizard — Network Tab, Routes for Traffic Page

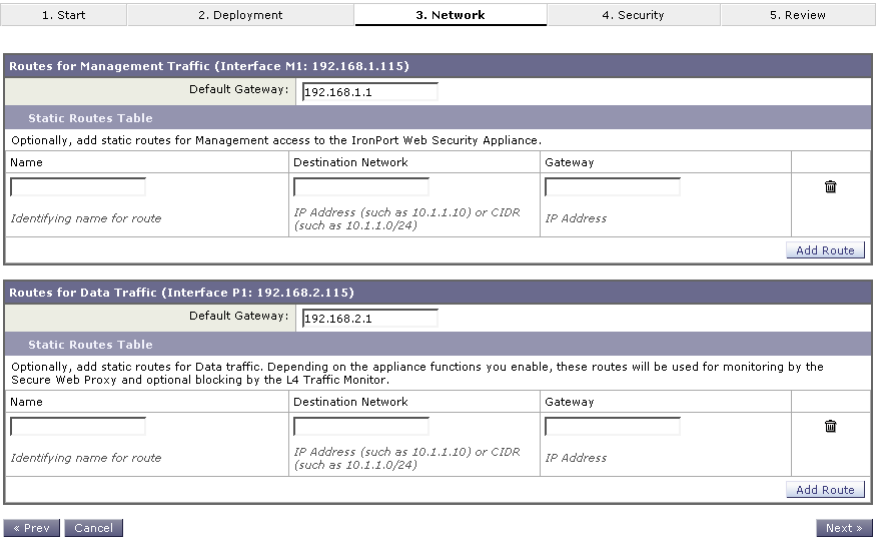

6. Configure the Routes for Management and Data Traffic options.

   The number of sections on this page depend on how you configured the "Restrict M1 port to appliance management services only" check box on the previous wizard page:

   - **Enabled.** When you use the Management interface for management traffic only, then this page includes two sections to enter gateway and static route table information, one for management traffic and one for data traffic. AsyncOS uses the management route information for management and data traffic, and data route information for data traffic.
   - **Disabled.** When you use the Management interface for both management and data traffic only, then this page includes one section to enter gateway and static route table information. AsyncOS uses the route information for both management and data traffic.

   Table 3-7 describes the Routes for Management and Data Traffic options.

Table 3-7 Routes for Management and Data Traffic Options in System Setup Wizard

Option	Description
Default Gateway	Enter the default gateway IP address to use for the traffic through the Management and/or Data interface.

Table 3-7 Routes for Management and Data Traffic Options in System Setup Wizard (Continued)

Option	Description
Static Routes Table	Optionally, you can add one or more static routes for management or data traffic. To add a static route, enter a name for the route, its destination network, and gateway IP address, and then click **Add Route**. A route gateway must reside on the same subnet as the Management or Data interface on which it is configured. To delete a static route you entered, click the Delete button next to the static route entry in the table. For more information about static routes, see "Configuring TCP/IP Traffic Routes" on page 396.

7. Click **Next**.

   When you configure the Web Security appliance as a web proxy in transparent mode and non-inline, you must connect it to a Layer 4 switch or a version 2 WCCP router. The Switch or Router Settings page appears.

   When you configure the Web Security appliance as a web proxy in forward mode, the Switch or Router Settings page does not appear. Go to step 10 on page 49.

Figure 3-9 System Setup Wizard — Network Tab, Switch or Router Settings Page

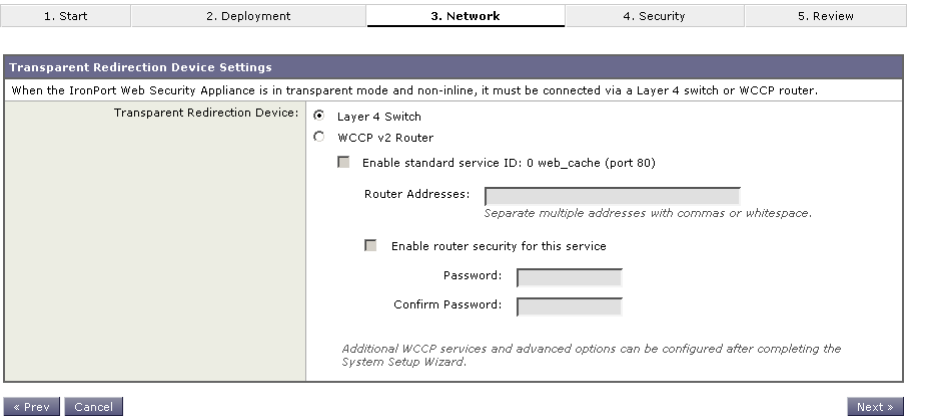

8. Specify whether the appliance is connected to a Layer 4 switch or a version 2 WCCP router.

   Note — If you connect the appliance to a version 2 WCCP router, you must configure the Web Security appliance to create WCCP services after you run the System Setup Wizard. For more information about creating WCCP services, see "Adding and Editing a WCCP Service" on page 405.

9. Click **Next**.

   The Administrative Settings page appears.

   Figure 3-10 System Setup Wizard — Network Tab, Administrative Settings Page

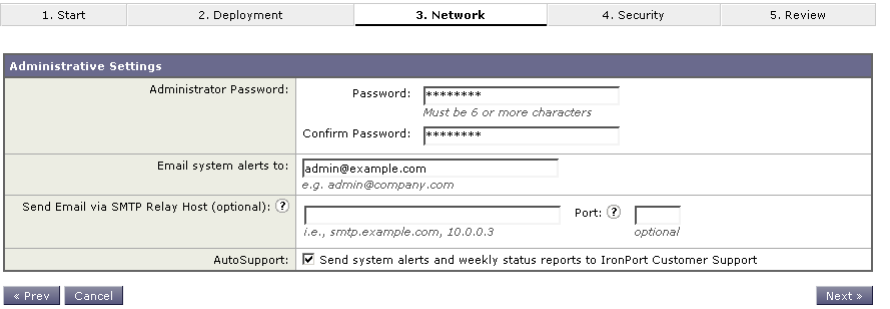

10. Configure the Administrative Settings options.

    Table 3-8 describes the Administrative Settings.

    Table 3-8 Administrative Settings in System Setup Wizard

Option	Description
Administrator Password	Enter a password to access the Web Security appliance. The password must be six characters or more.
Email System Alerts To	Enter an email address for the account to which the appliance sends alerts. For more information about alerts, see “Managing Alerts” on page 379.
Send Email via SMTP Relay Host	You can enter a host name or address for an SMTP relay host that AsyncOS uses for sending system generated email messages. Optionally, you can enter the port number, too. If no port number is defined, AsyncOS uses port 25. If no SMTP relay host is defined, AsyncOS uses the mail servers listed in the MX record. For more information about configuring the SMTP relay hosts, see “Configuring SMTP Relay Hosts” on page 386.
AutoSupport	Choose whether or not the appliance sends system alerts and weekly status report to IronPort Customer Support.

11. Click **Next**.

    The Security tab appears.

Step 4. Security

On the Security tab, you can configure which security services to enable, such as whether to block or monitor certain components. The Security tab contains one page.

1. Verify that you are viewing the Security tab.

Figure 3-11 System Setup Wizard — Security Tab

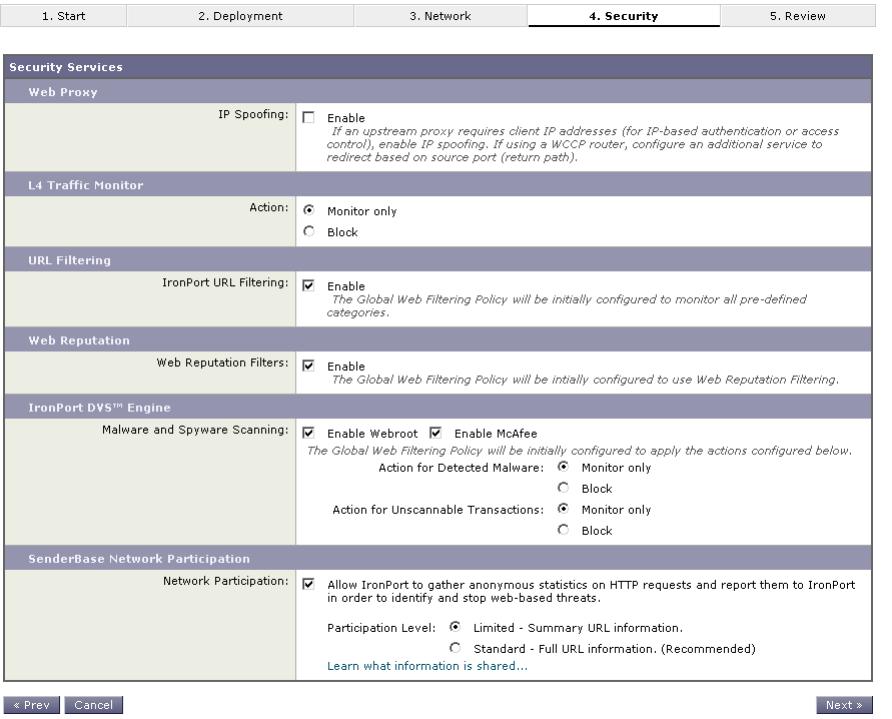

2. Choose the Security Services options.

Table 3-9 describes the Security options.

Table 3-9 Security Options in System Setup Wizard

Option	Description
Web Proxy	Choose whether or not to enable IP spoofing. Enable this option when a proxy server upstream in the network requires client IP addresses. This option only appears when on the Network Context page you specify that the network has an existing proxy. If you connect the appliance to a WCCP v2 router, you must create at least one WCCP service that redirects traffic based on the source port. For more information about creating WCCP services, see "Adding and Editing a WCCP Service" on page 405.
L4 Traffic Monitor	Choose whether the Layer-4 Traffic Monitor should monitor or block level 4 traffic. The L4 Traffic Monitor detects rogue traffic across all network ports and stops malware attempts to bypass port 80. You might choose to monitor traffic when you evaluate the Web Security appliance, and block traffic when you purchase and use the appliance. For more information, see "Configuring the L4 Traffic Monitor" on page 299.
URL Filtering	Choose whether or not to enable URL filtering. IronPort URL Filters allow you to control user access based on the category of a particular HTTP request. Enable this option when you want to restrict users from accessing particular types of websites. For more information, see "URL Filters" on page 207.
Web Reputation	Choose whether or not to enable Web Reputation filtering for the Global Policy Group. When you create custom access policy groups, you can choose whether or not to enable Web Reputation filtering. IronPort Web Reputation Filters is a security feature that analyzes web server behavior and assigns a reputation score to a URL to determine the likelihood that it contains URL-based malware. Enable this option when you want to identify suspicious activity and stop malware attacks before they occur. For more information, see "Web Reputation Filters" on page 227.

Table 3-9 Security Options in System Setup Wizard (Continued)

Option	Description
IronPort DVS Engine	Choose whether or not to enable malware and spyware scanning using Webroot or McAfee. If enabled, also choose whether to monitor or block detected malware, and whether to monitor or block unscannable transactions. You might choose to monitor malware and/or unscannable transactions when you evaluate the Web Security appliance, and block them when you purchase and use the appliance. You can further configure malware scanning after you finish the System Setup Wizard. For details, see "Configuring Anti-Malware Scanning" on page 246.
SenderBase Network Participation	Choose whether or not to participate in the SenderBase Network. If you participate, you can configure limited or full participation. The SenderBase Network is a threat management database that tracks millions of domains around the world and maintains a global watch list for Internet traffic. When you enable SenderBase Network Participation, the Web Security appliance sends anonymous statistics about HTTP requests to IronPort to increase the value of SenderBase Network data. For more information about the SenderBase Network, see "The SenderBase Network" on page 4.

3. Click Next.

   The Review tab appears.

Step 5. Review

The last tab of the System Setup Wizard displays a summary of the configuration information you chose. You can edit any of the configuration options, such as Network Settings or Deployment, by clicking the **Edit** button for each section.

1. Verify that you are viewing the Review tab.

Figure 3-12 System Setup Wizard — Review Tab

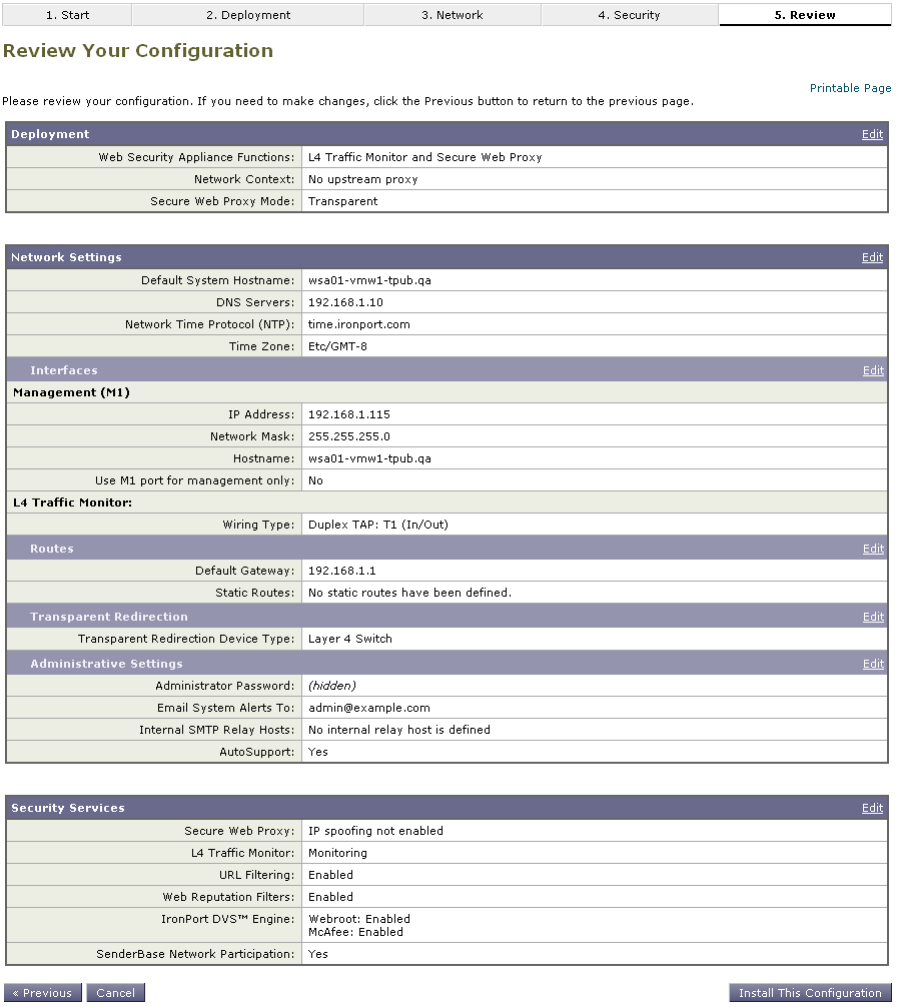

2. Review the configuration information. If you need to change an option, click the **Edit** button for that section.
3. Click **Install This Configuration** after you confirm the configuration is correct.

   The Web Security appliance applies the configuration options you selected.

   If you changed the Management interface IP address from the current value, then clicking **Install This Configuration** will cause the connection to the current URL to be lost. However, your browser will redirect itself to the new IP address. If you did not change the

IP address from the current value, the System Administration > System Setup > Next Steps page appears.

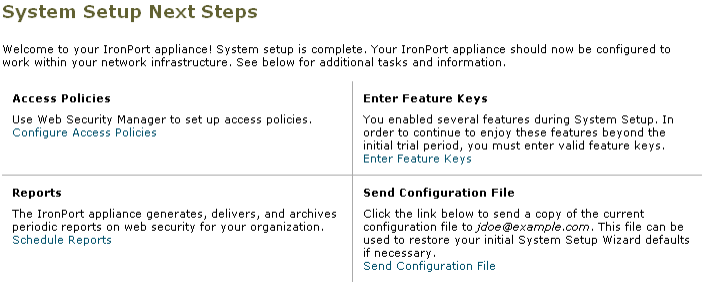

CHAPTER 4

Web Proxy Services

This chapter contains the following information:

ABOUT WEB PROXY SERVICES

A web proxy is a computer system or software that handles World Wide Web requests of clients by making requests of other servers on the web. The Web Security appliance can act as a web proxy if you enable the Web Proxy feature.

The Web Proxy service monitors and controls traffic that originates from clients on the internal network. Typically, the Web Proxy-enabled Web Security appliance is deployed between clients and the firewall where it intercepts requests for content from clients to servers.

You can configure the Web Proxy as one of the following types:

- **Transparent Proxy.** When the appliance is configured as a transparent proxy, clients are unaware of the Web Proxy. Client applications, such as web browsers, do not have to be configured to accommodate the appliance. You might want to configure the appliance as a transparent proxy because it eliminates the possibility of users reconfiguring their web browsers to bypass the appliance without knowledge of the administrator. To configure the appliance as a transparent proxy, you must connect it to an L4 switch or a WCCP router.

  For information about how to configure the appliance when you configure the proxy in transparent mode, see "Configuring Transparent Redirection" on page 402.

- **Explicit Forward Proxy.** In an explicit forward proxy configuration, the appliance acts on behalf of client web browsers to handle requests for servers on the web. Users must configure their web browsers to point to a single Web Security appliance. You might want to configure the appliance as an explicit forward proxy if you do not have an L4 switch or a WCCP router.

You can use the Web Security appliance in a network that includes another proxy server. For more information about how to deploy and configure the appliance when the network contains another proxy, see "Using the Web Security Appliance in an Existing Proxy Environment" on page 26.

Web Proxy Cache

By default, AsyncOS uses a web proxy cache to increase performance for users accessing the web in some cases.

You can edit the web proxy and proxy cache in the following ways:

- **Remove a URL from the cache.** Use the `evict` subcommand of the `webcache` CLI command to remove one or more URLs from the cache.
- **Specify a domain or URL to never cache.** Use the `ignore` subcommand of the `webcache` CLI command to specify one or more domains or URLs that the web proxy should never store in the proxy cache. You can included embedded regular expression (regex) characters in the URL you specify to never cache.

Each access log file entry includes transaction result codes that describe how the appliance resolved client requests. Transaction result codes indicate whether the transaction was served from the proxy cache or from the destination server. For more information about transaction result codes, see "Transaction Result Codes" on page 344.

CONFIGURING THE WEB PROXY

Web Proxy settings are configured as part of an initial setup using the System Setup Wizard. To enable Web Proxy services or modify proxy settings after an initial configuration, use the Security Services > Proxy Settings page. This page allows you to configure basic and advanced settings to customize proxy services.

Note — You must run the System Setup Wizard to change an initial transparent proxy or forward proxy configuration. All other settings can be modified using options on the Security Services > Proxy Settings page.

To edit the Web Proxy settings:

1. Navigate to the Security Services > Proxy Settings page.
2. Click **Edit Settings**.

Figure 4-1 Editing Web Proxy Settings

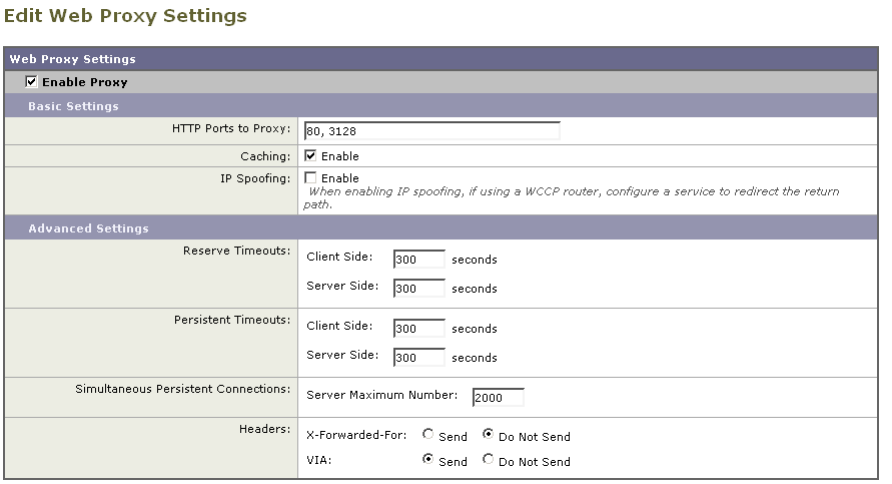

3. Verify the Enable Proxy field is selected.
4. Configure the basic and advanced Web Proxy settings defined in Table 4-1.

Table 4-1 Web Proxy Settings

Property	Description
HTTP Ports to Proxy	Enter which ports the Web Proxy monitors for HTTP requests. Default is 80 and 3128.

Table 4-1 Web Proxy Settings (Continued)

Property	Description
Caching	Choose whether or not the Web Proxy should cache requests and responses. Default is enabled.
IP Spoofing	Choose whether or not the Web Proxy should spoof IP addresses when forwarding requests to upstream proxies and servers. When IP spoofing is enabled, requests originating from a client retain the client's source address and appear to originate from the client rather than from the Web Security appliance. The appliance supports IP spoofing for transparent proxy configurations only. Note: When IP spoofing is enabled and the appliance is connected to a a WCCP router, configure a WCCP service to redirect the return path.
Reserve Timeouts	Enter how long the Web Proxy waits for more data from an idle client or server when the current transaction has not been completed. For example, if a client opens a connection and sends only half of the request, the Web Proxy waits for the amount of time specified for the client side reserve timeout for the rest of the request before closing the open connection. • **Client side.** The maximum number of seconds the Web Proxy keeps a connection open with an idle client. • **Server side.** The maximum number of seconds the Web Proxy keeps a connection open with an idle destination server. Default is 300 seconds for both client and server side reserve timeouts.

Table 4-1 Web Proxy Settings (Continued)

Property	Description
Persistent Timeouts	Enter how long the Web Proxy keeps open a connection to a client or server after a transaction has been completed. Keeping a connection open allows the Web Proxy to use it again for another request. For example, after a client finishes a transaction with google.com, the Web Proxy keeps the connection to the server google.com open for the amount of time specified in the server side persistent timeout if no other client makes a request for google.com. • **Client side.** The maximum number of seconds the Web Proxy keeps a connection open with a client on the network with no activity from the client. • **Server side.** The maximum number of seconds the Web Proxy keeps a connection open with a destination server with no activity from any client on the network to that server. Default is 300 seconds for both client and server side persistent timeouts. You might want to increase the server side persistent timeout if clients on the network frequently connect to the same server, or if the network has a relatively slow connection to outside servers. IronPort recommends keeping the default values. However, you might want to increase or decrease these values to keep connections open longer to reduce overhead used to open and close connections repeatedly. However, consider that if you increase the persistent timeout values, you also reduce the ability of the Web Proxy to open new connections if the maximum number of simultaneous persistent connections has been reached.
Simultaneous Persistent Connections (Server Maximum Number)	Enter the maximum number of connections (sockets) the Web Proxy keeps open with servers.
Headers	• **X-Forwarded-For.** Choose whether or not to forward HTTP "X-Forwarded-For" headers. Default is Do Not Send. Note: If the network contains an explicit forward upstream proxy that manages user authentication or access control using proxy authentication, you must enable the X-Forwarded-For header to send the client host header to the upstream proxy. • **VIA.** Choose whether or not to forward HTTP "VIA" headers. Default is Send.

5. Submit and commit your changes.

BYPASSING THE WEB PROXY

You can configure the Web Security appliance so client requests to or from particular addresses bypass all processing by the Web Proxy. The proxy bypass list only works for requests that are transparently redirected to the Web Proxy using an L4 switch or a WCCP v2 router. When the appliance is deployed in explicit forward mode, or when a client makes an explicit request to the Web Proxy, the request is processed by the Web Proxy.

You might want to create a proxy bypass list to accomplish any of the following:

- Prevent the Web Proxy from interfering with non-HTTP-compliant (or proprietary) protocols using HTTP ports that do not work properly when they connect to a proxy server.
- Ensure that traffic from a particular machine inside the network, such as a malware test machine, bypasses the Web Proxy and all its built-in security protection.

Define the proxy bypass list on the Web Security Manager > Proxy Bypass page.

Figure 4-2 shows a sample proxy bypass list.

Figure 4-2 Proxy Bypass List

To include an address in the proxy bypass list, click **Edit Settings**. You can enter multiple addresses separated by line breaks or commas. You can enter addresses using any of the following formats:

- IP address, such as 10.1.1.0
- CIDR address, such as 10.1.1.0/24
- Host name, such as crm.example.com
- domain names, such as example.com

Note — For the proxy bypass list to work with domain names, you need to connect the T1 and T2 network interfaces to the network *even if you do not enable the L4 Traffic Monitor*. For more information, see "How the Proxy Bypass List Works" on page 62.

When transactions bypass the Web Proxy, AsyncOS for Web records them in the proxy bypass logs. For more information about logging, see "Working with Log Subscriptions" on page 336.

Note — If the proxy bypass list contains an address that is a known malware address according to the L4 Traffic Monitor and the L4 Traffic Monitor sees a request for that address, then the request will still be blocked by the L4 Traffic Monitor. If you want to ensure traffic to that address is always allowed, you must also bypass the address from the L4 Traffic Monitor. For more information, see "How the L4 Traffic Monitor Works" on page 297.

How the Proxy Bypass List Works

When the Web Proxy receives an HTTP or HTTPS request, it checks both the source and destination IP address to see if it is in the proxy bypass list. If it is, the packet is sent to the next hop on the network. (In some cases, the packet is sent back to the transparent redirection device that redirected the packet, if the packet arrived on a WCCP service using GRE.)

The proxy bypass list works by matching the IP addresses of the request to an IP address in the proxy bypass list. When names are entered in the bypass list, the Web Proxy must resolve them to an IP address using DNS. The Web Proxy DNS resolves host names differently than domain names:

- **Host names.** Host names are resolved to IP addresses using DNS queries immediately after they are entered into the proxy bypass list. (An example host name is www.example.com.)
- **Domain names.** Domain names cannot be resolved to IP addresses using DNS queries, so the Web Proxy uses DNS snooping using the T1 and T2 network interfaces. (An example domain name is example.com, and it matches both www.example.com and webmail.example.com.)

Because of these differences, if the proxy bypass list contains only IP addresses and host names, then the Web Proxy can easily match the IP address in the request header to the IP addresses in the proxy bypass list.

However, for the proxy bypass list to work with domain names, you must connect both the T1 and T2 network interfaces (if using simplex mode) or just connect the T1 network interface (if using duplex mode) to the network *even if you do not enable the L4 Traffic Monitor*. However, the proxy bypass list only bypasses the Web Proxy scanning. It does not bypass the L4 Traffic Monitor.

Note — If the transparent redirection device is a WCCP router, some are intelligent enough to not forward any other packets to the Web Proxy for the same session. In this case, the packets are not physically sent to the Web Proxy for the rest of the session and are truly bypassing it for the rest of the session.

Using WCCP with the Proxy Bypass List

When the Web Security appliance is configured to use a WCCP v2 router, you must ensure that all WCCP services defined in the Web Security appliance use the same forwarding and return method (either L2 or GRE) to work properly with the proxy bypass list. If the forwarding and return methods do not match, some WCCP enabled routers will act inconsistently.

For more information, see “Working with the Forwarding and Return Method” on page 404.

PROXY USAGE AGREEMENT

You can configure the Web Security appliance to inform users that it is filtering and monitoring their web activity. The appliance does this by displaying an end-user acknowledgement page when a user first accesses a browser after a certain period of time. When the end-user acknowledgement page appears, users must click a link to access the original site requested or any other website. For more information about end-user acknowledgement pages, see "End-User Acknowledgement Page" on page 194.

CONFIGURING CLIENT APPLICATIONS TO USE THE WEB PROXY

Web browsers and other user agents sometimes need to know how to connect to the Web Proxy in order to access the World Wide Web. When you deploy the Web Security appliance in explicit forward mode, you *must* configure client applications so they use the Web Proxy. If you deploy the appliance in transparent mode, you can *choose* whether or not to configure client applications to explicitly use the Web Proxy.

You can configure client applications to explicitly use the Web Proxy by using any of the following configuration methods:

- **Manual.** Manual configuration involves typing the Web Security appliance host name and port number, such as 3128, in each client application. If the appliance changes, you must edit each application individually. You might want to manually configure an application when you are testing proxy access on a single client machine. IronPort does not recommend manually configuring each client application to use the appliance Web Proxy.
- **Proxy auto-config (PAC) file.** For web browsers, you can configure each browser to use a PAC file to find the Web Proxy. Then you can edit the PAC file to specify the appliance Web Proxy information. For more information, see "Working with PAC Files" on page 65.

For more information about how to configure client applications to use a proxy, see the client application documentation.

WORKING WITH PAC FILES

A proxy auto-config (PAC) file is a text file that defines how web browsers can automatically choose the appropriate proxy server for fetching a given URL.

When you use a PAC file, you only need to configure each browser once with the PAC file information. Then, you can edit the PAC file multiple times to add, delete, or change Web Proxy connection information without editing each browser. This way you can configure the proxy information about your network in a centralized location and update it easily.

Note — Once a browser has read a PAC file, it stores it in memory for the remainder of the browser session.

You might want to use a PAC file for the following reasons:

- **Centralized management.** You can manage the PAC file in a single, central location.
- **Complex network environment.** If the network of proxy servers is complicated, you can create a PAC file to accommodate different server and client needs.
- **Changing network environment.** If your network environment is likely to change in the future, you can easily add, edit, or delete proxy servers in the PAC and have the changes automatically affect all browsers.
- **Failover.** If you have multiple proxy servers, you can provide redundancy in case of failure. You can either program the PAC file to be redundant, or if a failure occurs, change the PAC file to use a different proxy server.

  Note — Different browsers take different amounts of time to fail over to a secondary proxy. For example, Internet Explorer takes about 25 seconds, and Firefox takes about 50 seconds.

- **Load balancing.** If you have multiple proxy servers, you can use the PAC file to specify which requests go to which proxy server. For example, you might want users on one subnet to use a particular proxy and users on a different subnet to use a different proxy.

PAC File Format

The PAC file must include at least one JavaScript function, FindProxyForURL(url, host). The JavaScript function determines the appropriate proxy to use for each URL.

For example, if the Web Security appliance host name is WSA.example.com, you could create a PAC file that includes the following text:

```
function FindProxyForURL(url, host) { return “PROXY
WSA.example.com:3128; DIRECT”; }
```

Note — The port you specify in the FindProxyForURL() function should be a proxy port for the Web Security appliance configured on the Security Services > Proxy Settings page.

However, you can make PAC files more complex. For example, you can create a PAC file that instructs the browser to connect directly to the website under certain conditions, such as matching on a particular host name or IP address, and to use the proxy server in all other cases. You can create a PAC file that instructs applications to go directly to the website for servers on your intranet.

For more information about creating and using PAC files, see the following locations:

- http://en.wikipedia.org/wiki/Proxy_auto-config
- http://www.mozilla.org/catalog/end-user/customizing/enduserPAC.html
- http://wp.netscape.com/eng/mozilla/2.0/relnotes/demo/proxy-live.html

Note — Common convention is to use the .pac file extension for PAC file names.

Creating a PAC File for Remote Users

Some laptop users connect to the Internet both from inside your organization's network and outside the network. For these users, you can create a PAC file that informs the browser to connect to the Web Proxy when they are on the network, and to connect directly to web servers when they are not on the network.

To do this, make sure the PAC file is hosted on a web server that is DNS resolvable inside the network, but not DNS resolvable outside the network. This works because when you enter a URL for the PAC file location, the browser will always try to use the PAC file in the configured location. If the browser cannot resolve the URL, such as when it is outside the network, it tries to access all web sites directly instead. Then when the laptop connects to the network again, the browser can access the PAC file and will use the Web Proxy to access web sites.

Specifying the PAC File in Browsers

To use a PAC file, you must publish the PAC file in a location that can be accessed by each browser that needs to access it. When you configure a browser to use a PAC file, you can use either of the following methods:

- **Enter the PAC file location.** See "Entering the PAC File Location" on page 66.
- **Detect the PAC file location automatically.** See "Detecting the PAC File Location Automatically" on page 67.

Entering the PAC File Location

You can configure a browser to use a PAC file by specifying the exact location of the file. You might want to enter the exact PAC file location for laptop users who might need to use different proxy servers depending on their current location.

You can place the PAC file in the following locations:

- **Local machine.** You can place the PAC file on each client machine and configure the browsers to use it. You might want to use a local PAC file to test a PAC file before deploying it to the entire organization. Enter the path in the browser configuration. The path you enter depends on the browser type.

- **Web server.** You can place the PAC file on a web server that each client machine can access. For example, you can place the PAC file on an Apache or Microsoft IIS web server. Enter the URL in the browser configuration.
- **Web Security appliance.** You can place the PAC file on the Web Security appliance. You might want to put the PAC file on the Web Security appliance to verify every client machine can access it within the network. Enter the URL in the browser configuration. If the URL does not specify the PAC file name, the appliance returns default.pac if it exists.

  For more information about uploading PAC files to the Web Security appliance, see "Adding PAC Files to the Web Security Appliance" on page 69.

Detecting the PAC File Location Automatically

If a browser supports the Web Proxy Autodiscovery Protocol (WPAD), you can configure it to automatically detect the PAC file location. WPAD is a protocol that allows the browser determine the location of the PAC file using DHCP and DNS lookups.

Before fetching its first page, a web browser configured to automatically detect the PAC file location tries to find the PAC file using DHCP or DNS. Therefore, to use WPAD, you must set up either a DHCP server or a DNS server to direct web browser requests to the PAC file on a network server. However, not all browsers support DHCP to find the PAC file using WPAD.

This section includes some general guidelines for using WPAD with DNS "A" records. For more detailed information, or for information about using WPAD with DHCP, see the following locations:

- http://en.wikipedia.org/wiki/Web_Proxy_Autodiscovery_Protocol
- http://www.wpad.com/draft-ietf-wrec-wpad-01.txt
- http://www.microsoft.com/technet/isa/2004/plan/automaticdiscovery.mspx

When you use WPAD with DNS, each domain on the network can only use one PAC file for all users on a domain because only domain name can uniquely identify a PAC file using DNS. For example, users on host1.accounting.example.com and host2.finance.example.com can use different PAC files.

To use WPAD with DNS:

1. Rename the PAC file to wpad.dat.
2. Create an internally resolvable DNS name that starts with "wpad," such as wpad.example.com.
3. Place wpad.dat in the root directory of the website that will host the file, such as wpad.example.com. For information about placing the file on the Web Security appliance, see "Uploading PAC Files to the Appliance" on page 69.

   Note — Due to a bug in Internet Explorer 6, create a copy of wpad.dat and change the file name to wpad.da to work with Internet Explorer 6 users. For more information, see http://www.microsoft.com/technet/isa/2004/ts_wpad.mspx.

4. Configure the web server to set up .dat files with the following MIME type:

   ```
   application/x-ns-proxy-autoconfig
   ```

   Note — If you place wpad.dat on the Web Security appliance, the appliance does this for you already.

ADDING PAC FILES TO THE WEB SECURITY APPLIANCE

You can configure browsers to explicitly use the Web Proxy by using proxy auto-config (PAC) files. When you use PAC files, you can place them on the Web Security appliance, and then configure the browsers by either entering the URL of a PAC file on the appliance or by configuring the browsers to automatically detect the PAC file using the Web Proxy Autodiscovery Protocol (WPAD).

You can add multiple PAC files to the appliance. You might want to add multiple PAC files if the appliance is used by multiple domains on the network. You can use one PAC file for all browsers on a domain.

When you add a PAC file to the appliance, you can specify one or more ports the appliance uses to listen for PAC file requests.

When a browser asks for a PAC file, the appliance sends back the file using HTTP. The PAC file is returned using MIME type `application/x-ns-proxy-autoconfig`.

Note — When browsers are configured to use a PAC file on the appliance, the URL should include the PAC file name. If the URL does not specify the PAC file name, the appliance uses default.pac if it exists and returns an error if it does not.

For more information about PAC files, see “Working with PAC Files” on page 65.

Uploading PAC Files to the Appliance

To store PAC files on the Web Security appliance:

1. Navigate to Security Services > Proxy Auto-Configuration File Hosting, and click **Enable and Edit Settings**.

   The Edit Proxy Auto-Configuration File Hosting Settings page appears.

   Figure 4-3 Editing the PAC File Host Settings

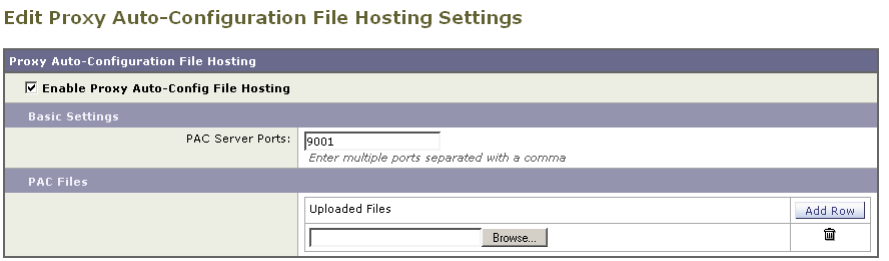

2. In the PAC Server Ports field, enter one or more port numbers the Web Security appliance should use to listen for PAC file requests.

   Note — Verify that the ports you use here are not listed as an HTTP port to proxy on the Security Services > Proxy Settings page. For example, if you want to use port 80 as the

PAC server port, you must first delete port 80 from the HTTP Ports to Proxy field if configured.

3. Click **Browse** to upload a PAC file from your local machine to the appliance.
4. Navigate to the PAC file location, select it, and click **Open**.
5. If you want to add another PAC file, click **Add Row**, and repeat steps 3 through 4.
6. Submit and commit your changes.

WPAD Compatibility with Netscape and Firefox

Netscape and Firefox browsers only use DNS to automatically detect PAC files using WPAD. Therefore, if you want Netscape and Firefox browsers to automatically detect a PAC file stored on the Web Security appliance, you must complete the following steps:

- Name the PAC file wpad.dat.
- Go to the Security Services > Proxy Settings page, and remove port 80 from the HTTP Ports to Proxy field.
- Use port 80 as the PAC Server Port when you upload the file to the appliance.

For more information about using WPAD, see "Detecting the PAC File Location Automatically" on page 67.

Note — The steps listed here also work with Internet Explorer, however, for Internet Explorer version 6, you should create a copy of wpad.dat and name it wpad.da.

ADVANCED PROXY CONFIGURATION

AsyncOS includes the `advancedproxyconfig` CLI command so you can configure more advanced Web Proxy configurations, such as authentication and DNS parameters.

The `advancedproxyconfig` command includes the following subcommands:

- **Authentication.** Configure authentication parameters, such as the number of outstanding concurrent Basic or NTLMSSP authentication requests to be authenticated by the authentication server and whether or not to log the username that appears in the request URI. You can also use the `authentication` subcommand to enable the user acknowledgment page. For more information about the user acknowledgment page, see "Proxy Usage Agreement" on page 63.
- **Caching.** Configure advanced Web Proxy caching options, such as:
  - Whether or not to ignore client requests to not retrieve content from the proxy cache
  - Whether or not to cache content from an untrusted server

  You can configure the parameters separately by selecting "Customized Mode," or you can choose a predefined set of parameter values. You can choose the following modes:
  - **Safe mode.** This mode uses less caching.
  - **Optimized mode.** This mode uses moderate caching.
  - **Aggressive mode.** This mode uses aggressive caching.
- **DNS.** Configure DNS-related options, such as the time to cache results of DNS errors and whether or not the Web Proxy should issue an HTTP 302 redirection on DNS lookup failure.
- **FTP.** Configure the login name and password to use for anonymous FTP access and whether or not to allow active mode for FTP transfers.
- **HTTPS.** Configure the logging style for URIs used in HTTPS transactions. You can choose to record the full URI ("fulluri") or just a portion of the URI with the query portion removed ("stripquery").
- **WCCP.** Configure the amount of logging detail to use to debug WCCP related issues.
- **Miscellaneous.** Configure whether or not the Web Proxy should respond to health checks from L4 switches and whether or not the Web Proxy should perform dynamic adjustment of TCP receive window sizes.

Each submenu command is discussed in the detail tables below. For the Default Value column, a string means a name or list of characters such as "hello world."

Authentication Options

Table 4-2 describes the authentication options for the `advancedproxyconfig` CLI command.

Table 4-2 advancedproxyconfig CLI Command—Authentication Options

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Would you like to forward Authorization request headers to a parent proxy?	Yes, No (Boolean)	No	Yes	This controls forwarding the 'Proxy-Authorization' in all deployments and in case of transparent mode with authentication enabled on the proxy, this also controls forwarding of the 'Authorization' header that the proxy receives during the authentication.
Enter the Proxy Authorization Realm to be displayed in the end user authentication dialog	String	"IronPort Web Security Appliance"	No	Proxy Authorization Realm displayed in the End User Authentication dialog.
Would you like to log the username that appears in the request URI?	Yes, No (Boolean)	No	No	If enabled, '<username>:xxxxx' is logged i.e the username is displayed and the password is represented as a string, 'xxxxx'. If disabled, both username and password are stripped. Note that the actual password is never displayed regardless of the value of this variable.
Would you like to turn on presentation of the User Acknowledgement page?	Yes, No (Boolean)	No	No	Enable or disable Acknowledgement page.
Enter maximum time to remember User Acknowledgement (in seconds):	30 - 2678400	86400	No	Maximum time to remember User Acknowledgement. From 30 seconds to one month (2678400).
Enter maximum idle timeout for User Acknowledgement based on IP Address (in seconds):	30 - 2678400	14400	No	Maximum idle timeout for User Acknowledgement based on IP Address. From 30 seconds to one month (2678400).

Table 4-2 advancedproxyconfig CLI Command—Authentication Options (Continued)

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Would you like to use Group Membership attribute while doing a directory lookup?	Yes, No (Boolean)	No	No	Choose whether or not AsyncOS should use the group membership attribute when doing a directory lookup. If you do not want to display empty authentication groups and fetch groups whose group membership attribute is different, choose Yes.
Would you like to enable referrals for LDAP?	Yes, No (Boolean)	No	Yes	Choose whether or not the Web Proxy should perform LDAP queries on a referred LDAP server. You might want to disable this option if a referred LDAP server is unavailable to the Web Security appliance.
Would you like to enable secure authentication?	Yes, No (Boolean)	No	Yes/No (Web Proxy restarts when it needs to listens on fewer or additional ports)	Choose whether or not the Web Proxy redirects clients to securely pass authentication credentials to the Web Proxy using HTTPS. For more information on this feature, see "Sending Authentication Credentials Securely" on page 279.
Would you like to use surrogates for explicit forward mode requests?	Yes, No (Boolean)	No	No	Choose whether or not to configure surrogate properties when the appliance is deployed in explicit forward mode even when secure authentication is not enabled. When you choose Yes, the CLI presents additional options you can configure. **Note:** This option only appears when you disable secure authentication.

Table 4-2 advancedproxyconfig CLI Command—Authentication Options (Continued)

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Enter the redirect port for secure authentication.	1 to 65535	443	Yes/No (Web Proxy restarts when it needs to listens on fewer or additional ports)	Enter the port to use for redirecting requests using HTTPS. IronPort recommends using a port greater than 1023. For more information on configuring this option, see "Configuring Global Authentication Settings" on page 271. **Note:** This option only appears when you enable secure authentication.
Enter surrogate type for authentication.	Cookie, IP	Cookie	No	This setting specifies the way that transactions used for authenticating the client are associated with a user (either by IP address or using a cookie) after the user has authenticated successfully. For more information on configuring this option, see "Configuring Global Authentication Settings" on page 271.
Enter the authentication cookie type.	Persistent, Session	Persistent	No	When you choose cookie as the surrogate type for authentication, you can choose either persistent or session cookies. For more information on configuring this option, see "Configuring Global Authentication Settings" on page 271.

Table 4-2 advancedproxyconfig CLI Command—Authentication Options (Continued)

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Enter the hostname to redirect clients for authentication.	String	Appliance host name	No	Enter the short host name of the network interface on which the Web Proxy listens for incoming connections. When you enable secure authentication, the Web Proxy uses this host name in the redirection URL sent to clients for authenticating users. For more information on configuring this option, see "Configuring Global Authentication Settings" on page 271.
Enter the surrogate timeout.	Time in seconds	3600	No	This setting specifies how long the surrogate (IP address or cookie) can be used before requiring authentication credentials again. For more information on configuring this option, see "Configuring Global Authentication Settings" on page 271.

Caching Options

The Caching submenu provides four options to set the advanced caching mode.

Table 4-3 describes the caching options for the Customized Mode option in the `advancedproxyconfig` CLI command.

Table 4-3 advancedproxyconfig CLI Command—Caching Options

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Would you like to allow objects with a heuristic expiration time to be served as not-modified If-Modified-Since hits from cache?	Yes, No (Boolean)	Yes	No	0 = favor freshness on IMS to objects with heuristic expiration time 1 = favor bandwidth conservation

Table 4-3 advancedproxyconfig CLI Command—Caching Options (Continued)

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Would you like to allow ETAG mismatch on client revalidations?	Yes, No (Boolean)	No	No	In some cases, the server might report different ETags for the same version of the same file. This can be seen, for example, with clustered IIS servers. In these cases, requiring both a last modified time (LMT) match and an ETag match on client revalidations would lead to a lot of misses, so it should be sufficient just to match the LMT if it is given. **Note:** Setting this to 1 is not HTTP-compliant.
Would you like to allow caching when requests are authenticated by the origin server?	Yes, No (Boolean)	No	Yes	Allow caching for requests authenticated by origin server.
Would you like to allow caching from servers whose DNS results do not match the TCP destination IP (not trustworthy and applicable only in transparent modes)?	Yes, No (Boolean)	No	Yes	Allow caching from servers whose DNS results do not match the TCP destination IP.
Enter the Heuristic maximum age to cache the document with Last-Modified Time but no actual caching value (in seconds):	Time in seconds	86400	No	Heuristic maximum age to cache the document with LMT but no actual caching value.
Enter the Heuristic maximum age to cache the document without Last-Modified Time and no actual caching value (in seconds):	Time in seconds	0	No	Heuristic maximum age to cache the document without LMT and no actual caching value.

Table 4-3 advancedproxyconfig CLI Command—Caching Options (Continued)

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Enter the Heuristic age to cache errors (HTTP_SERVICE_UNAVAIL, HTTP_GATEWAY_TIMEOUT etc) (in seconds):	Time in seconds	300	No	Heuristic age to cache errors (HTTP_SERVICE_UNAVAIL, HTTP_GATEWAY_TIMEOUT etc).
Would you like proxy to ignore client directive to not fetch content from the cache?	Yes, No (Boolean)	No	No	Disable/Enable ignoring of the client directive to not fetch content from the cache. Enabling this is not HTTP compliant.
Enter the time interval during which reload requests must be ignored by the proxy (in seconds):	Time in seconds	0	No	Disable/Enable reload requests to be ignored for the specified time interval. This allows reload requests to be ignored for a certain amount of time, even though it is not HTTP-compliant. You might want to enter a value greater than zero to improve bandwidth usage.
Would you like to allow proxy to convert reload requests into max-age requests?	Yes, No (Boolean)	No	No	Allow reload requests to be converted into max-age requests (not HTTP-compliant, but may improve bandwidth usage). This gets its max-age value from "ignoreReloadTime."
Time in seconds after which an explicit IMS Refresh request must be issued:	Time in seconds	300	No	Time in seconds after which an explicit IMS Refresh request must be issued.

DNS Options

Table 4-4 describes the DNS options for the `advancedproxyconfig` CLI command.

Table 4-4 advancedproxyconfig CLI Command—DNS Options

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Enter the time to cache successful DNS results if DNS does not provide TTL (in seconds):	0 - 86400	300	No	Time to cache successful DNS results if DNS does not provide TTL.
Enter the time to cache results of DNS errors (negative DNS caching) (in seconds):	0 - 86400	30	No	Set to 0 to be HTTP compliant.
Enter the URL format for the HTTP 302 redirection on DNS lookup failure:	String with EUN page variables	http://www.%H.com/%u	No	URL format for the HTTP 302 redirection on DNS lookup failure. See Table 10-2, "Variables for Customized End-User Notification Pages," on page 188 for the list of valid variables.
Would you like the proxy to issue a HTTP 302 redirection on DNS lookup failure?	Yes, No (Boolean)	Yes	Yes	Disable/Enable automatic HTTP 302 redirection on DNS lookup failure.
Would you like proxy not to automatically failover to DNS results when upstream proxy (peer) is unresponsive?	Yes, No (Boolean)	No	Yes	Disable/Enable automatic failover to DNS results when upstream proxy (peer) is unresponsive.
Find web server by: 0 = use DNS answers in order, 1 = use client supplied address then DNS, 2 = use ONLY client supplied address:	0, 1, 2	1	Yes	Specify how the appliance should find the location of the requested web server.

FTP Options

Table 4-5 describes the FTP options for the `advancedproxyconfig` CLI command.

Table 4-5 advancedproxyconfig CLI Command—FTP Options

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Enter the login name to be used for anonymous FTP access:	String	anonymous	No	Anonymous FTP login name.
Enter the password to be used for anonymous FTP access:	String	proxy@	No	Anonymous FTP login password.
Would you like to use active FTP transfer mode when passive mode fails?	Yes, No (Boolean)	No	No	Choose whether or not to allow FTP transfers to use active mode if passive mode fails.
Enter the range of port numbers for the proxy to listen on for active FTP connections.	port1-port2 (string) 1024 - 65535	10000 - 10001	No	When you enable active mode for FTP transfers, enter the range of ports the appliance can use for establishing a data connection. If a port is being used, the Web Proxy chooses the next port in the range until it finds an available port.

HTTPS Options

Table 4-6 describes the HTTPS options for the `advancedproxyconfig` CLI command.

Table 4-6 advancedproxyconfig CLI Command—HTTPS Options

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
HTTPS URI Logging Style:	fulluri or stripquery	fulluri	Yes	You can log the entire URI (fulluri), or a partial form of the URI with the query portion removed (stripquery). However, even when you choose to strip the query from the URI, personally identifiable information may still remain.

WCCP Options

Table 4-7 describes the WCCP options for the `advancedproxyconfig` CLI command.

Table 4-7 advancedproxyconfig CLI Command—WCCP Options

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Enter the log level for debugging WCCP:	0 - 10	0	Yes	WCCP log level

Miscellaneous Options

Table 4-8 describes the miscellaneous options for the `advancedproxyconfig` CLI command.

Table 4-8 advancedproxyconfig CLI Command—Miscellaneous Options

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Would you like proxy to respond to health checks from L4 switches (always enabled if WSA is in L4 transparent mode)?	Yes, No (Boolean)	No	Yes	Disable/Enable support for responding to health checks from L4 switches (always enabled if WSA is in L4 transparent mode). L4 switches issue 'HEAD / HTTP/1.0' requests directed at the proxy to ensure that it is responding.
Would you like proxy to perform dynamic adjustment of TCP receive window size?	Yes, No (Boolean)	Yes	Yes	Disable/Enable dynamic adjustment of TCP receive window size.
Enable custom EUN pages?	Yes, No (Boolean)	No	Yes	Choose whether or not to enable the ability to upload user-defined end-user notification pages to the appliance using FTP. For more information, see "Editing IronPort Notification Pages" on page 187.
Enable caching of HTTPS responses?	Yes, No (Boolean)	No	No	Choose whether or not the Web Security appliance should store HTTPS responses in the web cache.

Table 4-8 advancedproxyconfig CLI Command—Miscellaneous Options (Continued)

Option	Valid Values	Default Value	Must Restart Web Proxy?	Description
Enter minimum idle timeout for checking unresponsive upstream proxy (in seconds).	Time in seconds	10	No	The minimum amount of time the Web Proxy waits before checking if an upstream proxy is still unavailable.
Enter maximum idle timeout for checking unresponsive upstream proxy (in seconds).	Time in seconds	86400	No	The maximum amount of time the Web Proxy waits before checking if an upstream proxy is still unavailable.
Do you want proxy to listen on P2?	Yes, No (Boolean)	No	Yes	Choose whether or not the Web Proxy should listen for web requests on the P2 network interface.

CHAPTER 5

Working with Policies

This chapter contains the following information:

WORKING WITH POLICIES OVERVIEW

The Web Security appliance allows you to define policies to enforce your organization's acceptable use policies by controlling access to the Internet. You can create groups of users and apply different levels and types of access control to each group.

For example, you can configure the appliance to enforce the following types of policies:

- Users in the Marketing group can access a competitor's website, but other users cannot.
- Guest users on customer-facing machines, such as computers in a company store, cannot access banking sites, but employees can.
- No users can access gambling sites. Instead, when they try to view a gambling site, they see a web page that explains the organization's policies.
- All users trying to access a particular site that no longer exists are redirected to a different site.
- All users except those in IT are blocked from accessing potential malware sites, but users in IT can access them for testing purposes, and the downloaded content is scanned for harmful objects.
- All requests for streaming media are blocked during business hours, but allowed outside of business hours.
- All requests from a particular user agent, such as a software update program, are allowed without requiring authentication.

To enforce organizational policies, you define different policies in the Web Security appliance. The appliance uses different types of policies for different functions. For more information about the types of policies, see "Policy Types" on page 85.

When you work with policies, you create policy groups. After you create policy groups, you can define the access control settings for each group. For more information about working with policy groups, see "Working with Policy Groups" on page 87.

After you have created policies, you can figure out which policy groups apply to a particular client request for troubleshooting purposes. For example, you can find out if user jsmith tries to open a Firefox browser to the URL http://www.google.com, then which policy groups apply to the transaction. For more information about tracing policies, see "Tracing Policies" on page 98.

Note — The Web Security appliance is permissive by default. That is, requests are allowed unless specifically blocked in a policy group.

POLICY TYPES

The Web Security appliance uses multiple types of policies to enforce organizational policies and requirements.

- **Identities.** "Who are you?"
- **Decryption policies.** "To decrypt or not to decrypt?"
- **Routing policies.** "From where to fetch content?"
- **Access policies.** "To allow or block the transaction?"

You use the policies together to create the behavior you need or expect when clients access the web.

To define policies, you create policy groups. After you create policy groups, you can define the access control settings for each group. For more information about working with policy groups, see "Working with Policy Groups" on page 87.

All policy types have a global policy group that maintains default settings and rules that apply to web transactions not covered by another policy. For more information on global policies, see "Working with Policy Groups" on page 87.

Identities

An identity is a policy that identifies the user making a request. This is the only policy where you can define whether or not authentication is required. An identity addresses the question, "who are you?" However, identities do *not* specify a list of users who are *authorized* to access the web. You specify authorized users in the other policy types after you specify the identity to use.

All other policies you create must specify an identity.

Configure identities on the Web Security Manager > Identities page. For more information about identities, see "Identities" on page 103.

Decryption Policies

A decryption policy determines whether or not an HTTPS connection should be decrypted, passed through, or dropped. It addresses the question, "to decrypt or not to decrypt?"

The appliance uses decryption policies to evaluate HTTPS requests. The decryption policy group that applies to an HTTPS request determines whether the appliance drops the connection, passes it through without decryption, or decrypts the connection and subsequently evaluate the decrypted request and response against the defined access policy groups.

Configure decryption policy groups on the Web Security Manager > Decryption Policies page. For more information about decryption policy groups, see "Decryption Policies" on page 149.

Routing Policies

A routing policy determines to where to pass the client request, either to another proxy or to the destination server. It addresses the question, "from where to fetch content?"

You can use this policy type to select a group of upstream proxies configured for load balancing or failover.

Configure routing policies on the Web Security Manager > Routing Policies page. For more information about routing policies, see "Working with External Proxies" on page 137.

Access Policies

An access policy determines whether to allow or block HTTP and decrypted HTTPS transactions. It addresses the question, "to allow or block the transaction?"

Access policies determine how the appliance controls access to services, applications, and objects on the web for HTTP and decrypted HTTPS requests. The appliance uses access policies to evaluate and scan HTTP requests and HTTPS requests designated for decryption.

Configure access policy groups on the Web Security Manager > Access Policies page. For more information about access policy groups, see "Access Policies" on page 119.

WORKING WITH POLICY GROUPS

A policy group is an administrator defined configuration that allows you to apply acceptable use policies to specific categories of users. After you create policy groups, you can define the access control settings for each group.

You can create as many user defined policy groups as required to enforce the proper access control. The Web Security appliance displays policy groups together in a policies table.

All policies have a default, global policy group that applies to a transaction if none of the user defined policy groups apply. A global policy group maintains default settings and rules that apply to web transactions not covered by another policy. This group appears in the last row of a policies table, and AsyncOS applies its rules last if no other matching occurs.

Creating Policy Groups

You can create policy groups based on combinations of several criteria, such as client subnet or the URL category of the destination site. You must define at least one criterion for policy group membership. When you define multiple criteria, the client request must meet all criteria to match the policy group.

Options used to configure policy groups allow you to specify exceptions to global policy settings and control access to services for groups of users.

For more information about creating policy groups for the different policy types, see the following locations:

- “Creating Identities” on page 112
- “Creating Access Policies” on page 124
- “Creating Decryption Policies” on page 172
- “Creating Routing Policies” on page 145

Using the Policies Tables

The policies table is an ordered list of policy groups and the settings you configure for each filtering component. It displays policy groups by row and access control components by column. The access control components you can define vary by policy type.

Figure 5-1 on page 88 shows the access policies table.

Figure 5-1 Access Policies Table

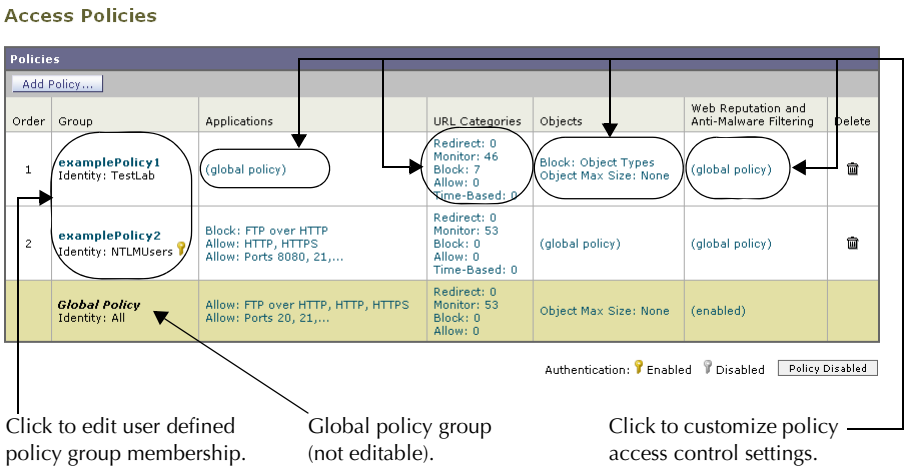

Figure 5-2 on page 88 shows the decryption policies table.

Figure 5-2 Decryption Policies Table

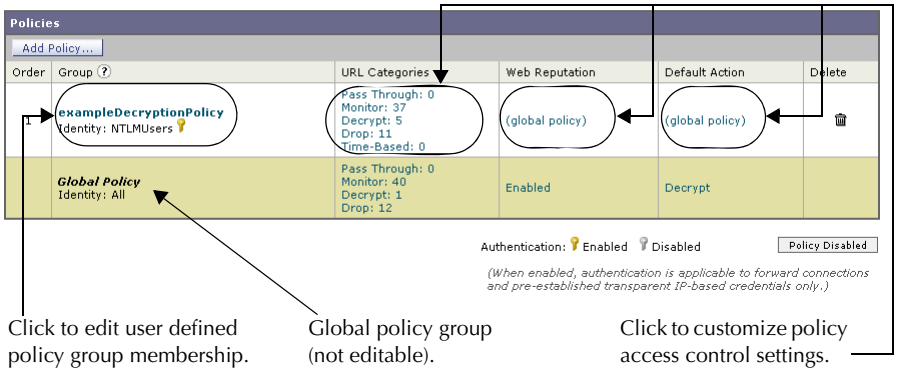

Any policy group that you create is added as a new row in the policies table. New policy groups inherit global policy settings for each access control component until you override them. To edit policy groups, click the links in each row.

When you create or configure a policy group, you define the following components:

- **Policy group membership.** Define how to group users that belong to the policy group. For user defined policy groups, you can group by different properties, such as client IP

address, authentication group or user name, or URL category. The properties you can define for a policy depends on the policy type.

Click the policy group name to edit the group membership requirements, such as client IP address and authentication requirements. A page is displayed where you can configure membership requirements.

Note — For global policies, you can only define the membership requirements for the global identity group and not for the global access, decryption, or routing groups. Global access, decryption, and routing groups always match all identities.

For more information about policy group membership, see "Policy Group Membership" on page 90.

- **Policy group access control settings.** Define how users in the group can use the Internet. The access control settings you can define depend on the policy type. For example, for routing policies, you define from which proxy group to fetch the content, and for access policies, you can use the Web Security appliance features, such as Web Reputation, anti-malware scanning, and more to determine whether or not to allow the client request.

  Click the link in the policy group row under the access control component you want to configure, such as URL Categories or Routing Destination. When you click a link in the table, a page is displayed where you can configure settings for that policy group.

  For more information on configuring control settings for each policy type, see the following sections:

  - "Controlling Access to HTTP Traffic" on page 128
  - "Controlling HTTPS Traffic" on page 176
  - "Creating Routing Policies" on page 145

POLICY GROUP MEMBERSHIP

All policy groups define which transactions apply to them. When a client sends a request to a server, AsyncOS receives the request, evaluates it, and determines to which policy group it belongs. AsyncOS applies the configured policy settings to a client request based on the client request's policy group membership.

Transactions belong to a policy group for each type of policy that is enabled. If a policy type has no user defined policy groups, then each transaction belongs to the global policy group for that policy type.

Policy group membership for a routing, decryption, and access policies is based on an identity and optional additional criteria. That means that *AsyncOS evaluates identity groups before the other policy types*. The Web Security appliance allows you to define some membership criteria at both the identity level and the non-identity policy level. However, any membership criteria you define at the non-identity level further narrows down the list of transactions that match the policy group.

Suppose you define an identity by subnet 10.1.1.0/24 and then create an access policy using that identity. The access policy membership applies to all IP addresses specified in the identity by default. You can then choose to configure the access policy membership so that it applies to a subset of the addresses defined in the identity, such as addresses 10.1.1.0-15.

For more information defining membership for each policy type, see the following sections:

- "Evaluating Identity Group Membership" on page 105
- "Evaluating Access Policy Group Membership" on page 122
- "Evaluating Decryption Policy Group Membership" on page 170
- "Evaluating Routing Policy Group Membership" on page 143

Authenticating Users versus Authorizing Users

The Web Security appliance separates where it authenticates users from where it authorizes users.

Authentication is the mechanism by which AsyncOS securely identifies a user. It answers the following questions:

- Who is the user?
- Is the user really whom he/she claims to be?

Authorization is the mechanism by which AsyncOS determines the level of access the user has to the World Wide Web. It answers the following questions:

- Is this user allowed to view this website?
- Is this user allowed to connect to this HTTPS server without the connection being decrypted?

- Is this user allowed to directly connect to the web server, or must it connect to another proxy server first?

AsyncOS can only authorize a user to access an Internet resource *after* it authenticates who the user is. AsyncOS authenticates users when it evaluates identity groups, and it authorizes users when it evaluates all other policy group types. What that means is the identity group indicates who is making the request, but does not indicate whether that client is allowed to make the request.

By separating authentication from authorization, you can create a single identity group that identifies a group of users and then you can create multiple policy groups that allow different levels of access to subsets of users in the group in the identity.

For example, you can create one identity group that covers all users in an authentication sequence. Then you can create an access policy group for each authentication realm in the sequence. You can also use this identity to create one decryption policy with the same level of access for all users in the identity.

Working with All Identities

You can create a policy group that specifies "All Identities" as the configured identity group. "All Identities" applies to every valid client request because by definition, every request either succeeds and has a user defined or global identity assigned to it or is terminated because it fails authentication.

When you create a policy group that uses All Identities, you must configure at least one advanced option to distinguish the policy group from the global policy group.

Typically, you use All Identities in policy while also configuring an advanced option, such as a particular user agent or destination (using a custom URL category). This allows you to create a single rule that makes an exception for a specific case instead of creating multiple rules to make the exception for the specific case. For example, you can create an access policy group whose membership applies to All Identities and a custom URL category for all intranet pages. Then you can configure the access policy control settings to disable anti-malware filtering and Web Reputation scoring.

Note — When a transaction matches a policy group that uses All Identities, the identity is not included in the access logs.

Policy Group Membership Rules and Guidelines

Consider the following rules and guidelines when defining policy group membership:

- AsyncOS evaluates identity groups before the other policy types.
- Advanced membership criteria for access, decryption, and routing policies further narrow down the list of transactions that match the policy group compared to the configured identity group.
- Define identity groups as broadly as possible. Then you can use the identity groups in other policy types and further narrow down membership as necessary.

- Define fewer, more generic decryption and routing policies as much as possible.
- If you need to define membership by URL category, only define it in the identity group when you need to exempt from authentication requests to that category. For other purposes, define membership by URL category in the access, decryption, or routing policy group. This can increase performance in most cases.

WORKING WITH TIME BASED POLICIES

The Web Security appliance provides the means to create time based policies by specifying time ranges, such as business hours, and using those time ranges to define access to the web. You can define policy group membership based on time ranges, and you can specify actions for URL filtering based on time ranges.

You might want to use time ranges to accomplish the following tasks:

- You can block access to high bandwidth sites, such as streaming media, or distracting sites, such as games, during business hours.
- You can route transactions to a particular external proxy after midnight when the other proxies are being serviced.
- You can allow larger files to be downloaded on the weekends.

Define time ranges on the Web Security Manager > Time Ranges page. You can create time ranges to define concepts such as "business hours" or "weekend shift." Then you can use the time ranges in the following locations:

- Policy group membership for a routing, access, or decryption policy.
- URL filtering settings for access policies.

When you define a time range, you can specify the day(s) of the week and the time of day. A transaction matches the time range when it occurs on one of the days specified and during the time specified. You can also multiple combinations of day and time in a single time range. For example, you can define a time range that applies to transactions that occur on Monday through Friday from 08:00 to 17:00 or on Saturday from 09:00 to 13:00.

Policies and URL filtering actions can be define inside or outside the defined time ranges.

Note — Because you can define time based policy group membership only for routing, access, and decryption policies, but not identities, you cannot create time based policies that define when users must authenticate. Authentication requirements are defined in identity groups, but time based policies are defined in other policy group types.

Creating Time Ranges

To create a time range:

1. Go to Web Security Manager > Time Ranges.
2. Click **Add Time Range**.

   The Add Time Range page appears.

Add Time Range

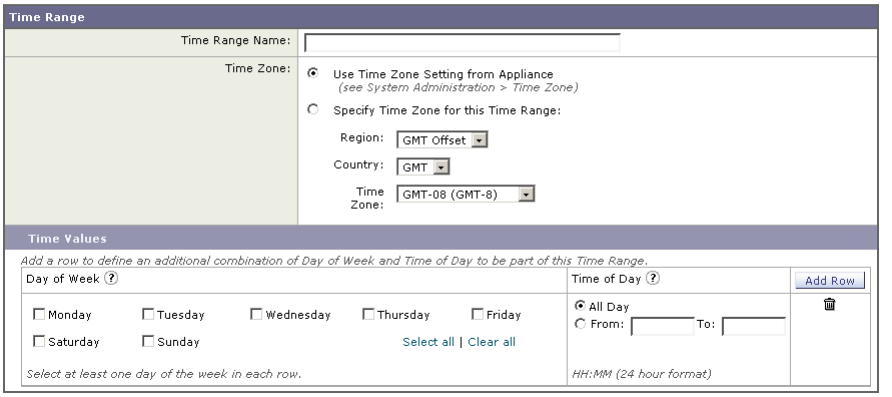

3. In the Time Range Name field, enter a name to use for the time range. Each time range name must be unique.
4. In the Time Zone section, choose whether to use the time zone setting on the Web Security appliance or a different time zone setting you configure.
5. In the Time Values section, define at least one row that specifies the days of the week and time of day to include in this time range.
    a. In the Day of the Week section, select at least one day.
    b. In the Time of Day section, choose All Day or enter a time range in the day using the From and To fields.

       Each time range includes the start time and excludes the end time. For example, entering 8:00 through 17:00 matches 8:00:00 through 16:59:59, but not 17:00:00.

       Midnight must be specified as 00:00 for a start time, and as 24:00 for an end time.

    Note — A transaction must occur on the day *and* in the time specified to match a row in the Time Values section. That means the Day of Week and Time of Day values have an "AND" relationship with each other within a single row.

6. Optionally, you can create additional time value rows by clicking **Add Row**.

    Note — When a time range includes multiple time value rows, a transaction can occur within any of the defined time values to match the time range. That means that multiple time value rows in a single time range have an "OR" relationship with each other.

7. Submit and commit your changes.

WORKING WITH USER AGENT BASED POLICIES

The Web Security appliance provides the means to create policies to define access to the web by the client application (user agent), such as a web browser, making the client request. You can define policy group membership based on user agents, and you can specify access control settings based on user agents.

You might want to specify user agents to accomplish the following tasks:

- You can exempt certain user agents from authentication. You might want to do this for client applications that cannot handle prompting users for authentication credentials. For more information about how to do this, see "Exempting User Agents from Authentication" on page 97.
- You can block access from particular user agents that you define.

You can configure user agents in the following locations:

- Policy group membership for all policy types, including identities.
- Application control settings for access policies.

Configuring User Agents for Policy Group Membership

When you define policy group membership for any policy type, you can expand the Advanced section to define membership by additional criteria, such as user agent. When you click the User Agents link, the Membership by User Agent page appears allowing you to define membership by user agent.

Figure 5-3 on page 96 shows the Membership by User Agent page for an identity policy group.

Figure 5-3 Defining Policy Group Membership by User Agent

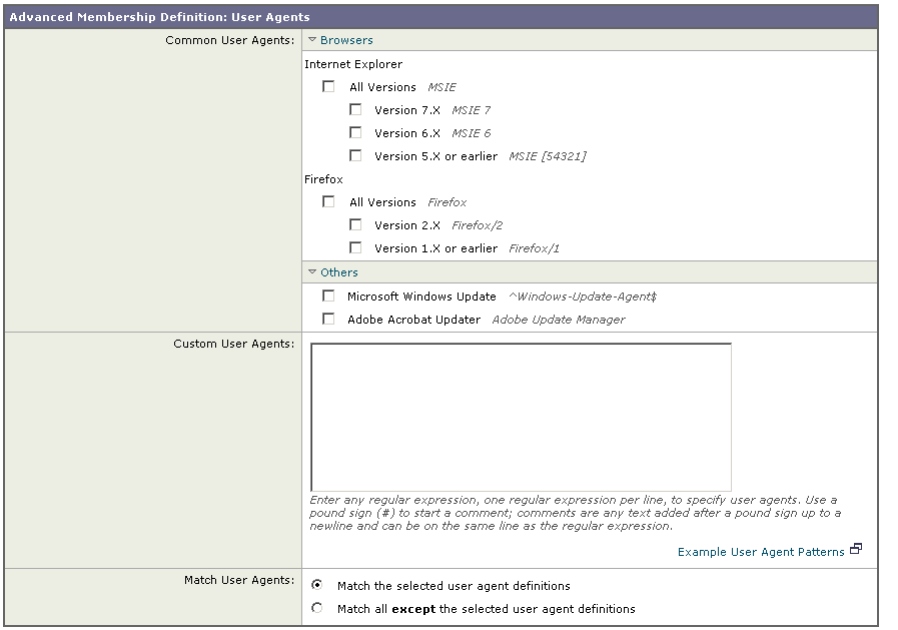

On this page, you can select as many user agents as desired. The web interface includes some of the more common user agents that you can select using a check box. You can also type a regular expression to define any user agent necessary.

For each user agent you select in the Common User Agents section, AsyncOS for Web creates a regular expression to define the user agent. However, if you select the All Versions option for each browser type, AsyncOS for Web creates a single regular expression that represents all versions of that browser instead an expression for each version. Creating one regular expression instead of multiple increases performance.

For example, when you select “Version 2.X” and “Version 1.X or earlier” for Firefox, AsyncOS for Web uses the following regular expressions:

```
Firefox/2
Firefox/1
```

However, when you select “All Versions” under Firefox, AsyncOS uses the following regular expression:

```
Firefox
```

Also, you can configure the policy group membership to either match the user agents you define, or matching all other user agents than the ones defined.

Exempting User Agents from Authentication

To exempt a user agent from authentication:

1. Create an identity policy group with membership that is based on the user agent to exempt.

   For more information about creating identities, see “Creating Identities” on page 112.

2. Do not require authentication for the identity policy group.
3. Place the identity policy group above all other identity policy groups that require authentication.
4. Submit and commit your changes.

TRACING POLICIES

The Web Security appliance web interface includes a tool that traces a particular client request and details how the Web Proxy processes the request. The Web Proxy evaluates the request against all committed policies and calculates other attributes, such as the web reputation score.

The policy trace tool allows administrators to troubleshoot when end users ask questions about Web Proxy behavior. It simulates client requests as if they were made by the end users and describes Web Proxy behavior. It can be a powerful troubleshooting or debugging tool, especially if you have combined many of the advanced features available on the Web Security appliance.

When you use the policy trace tool, the Web Proxy does not record the requests in the access log or reporting database.

Note — The policy trace tool explicitly makes requests even if the Web Security appliance is deployed in transparent mode.

You can trace policies on the System Administration > Policy Trace page.

To trace policies:

1. Navigate to the System Administration > Policy Trace page.

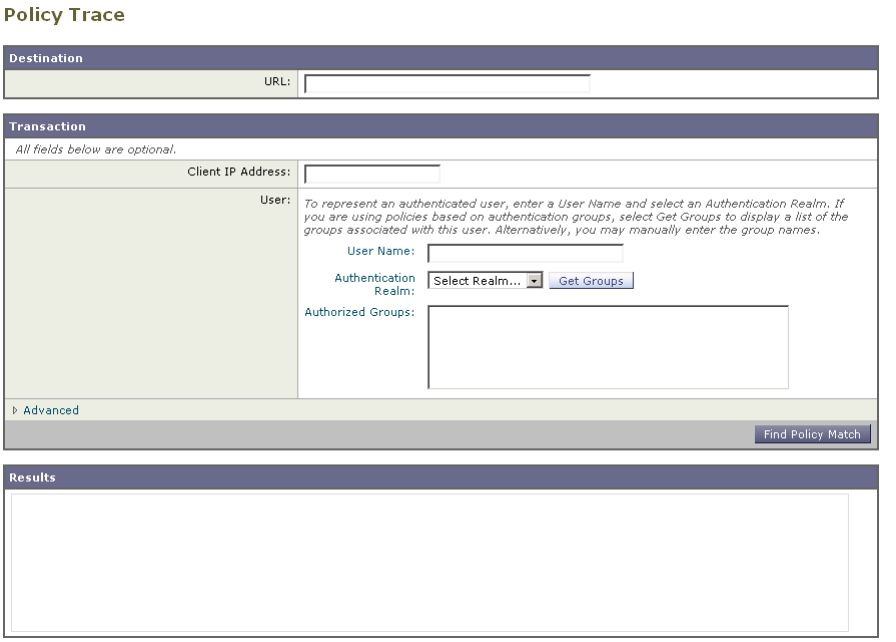

2. In the URL field, enter the URL in the client request to simulate.
3. Optionally, in the Client IP Address field, enter the IP address of the machine to simulate.

   Note — If no IP address is specified, AsyncOS uses localhost.

4. Optionally, you can simulate an authentication user by entering the following authentication requirements in the User area:
   - **User Name.** Enter the user name of the authentication user.
   - **Authentication Realm.** Choose an authentication realm.
   - **Authorized Groups.** If any of the policies use authentication groups, you can click **Get Groups** to get a list of all authentication groups this user is a member of on the selected realm from which to select.

   Note — For the authentication to work for the user you enter here, the user must have already successfully authenticated through the Web Security appliance.

5. Optionally, by expanding the Advanced section, you can configure additional settings to simulate a more specific user request that you want to trace.

   Figure 5-4 shows the expanded Advanced section.

   Figure 5-4 Policy Trace Feature Advanced Section

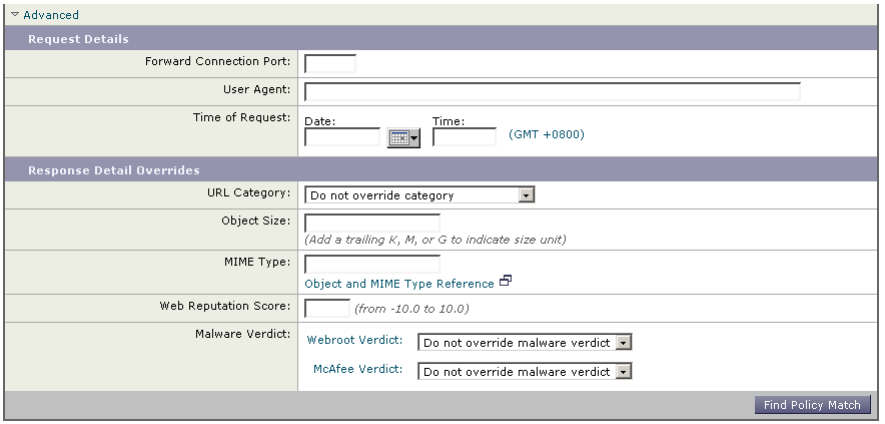

   The Advanced settings are divided into details of the transaction request to simulate and transaction response details to override.

6. Configure the transaction request information to simulate as desired. Table 5-1 describes the request side advanced settings you can configure.

Table 5-1 Policy Trace Advanced Settings for Requests

Setting	Description
Forward Connection Port	Select a specific proxy port to use for the trace request to test policy group membership based on proxy port.
User Agent	Specify the user agent to simulate in the request.
Time of Request	Specify the day of week and time of day to simulate in the request.

7. Configure the transaction response details to override as desired.

   You might want to override a transaction response detail to simulate how a different response value, such as a lower web reputation score, would affect the policies assigned to the transaction. Table 5-2 describes the response side advanced settings you can configure.

Table 5-2 Policy Trace Advanced Settings for Response Overrides

Setting	Description
URL Category	Choose whether or not to override the URL category of the transaction response.
Object Size	Enter the size of the response object in bytes. You can enter K, M, or G to represent Kilobytes, Megabytes, or Gigabytes.
MIME Type	Enter the MIME type.
Web Reputation Score	Enter the web reputation score from -10.0 to 10.0.
Malware Verdict	Choose whether or not to override the Webroot or McAfee scanning verdicts.

8. Click **Find Policy Match**.

   The policy trace tool displays the results in the Results area.

   Note — The **Find Policy Match** button turns into a **Cancel** button while the policy trace processes the parameters you enter. You can cancel the trace at any time.

Figure 5-5 on page 101 shows the Policy Trace page with some results from a policy trace.

Figure 5-5 Policy Trace Results

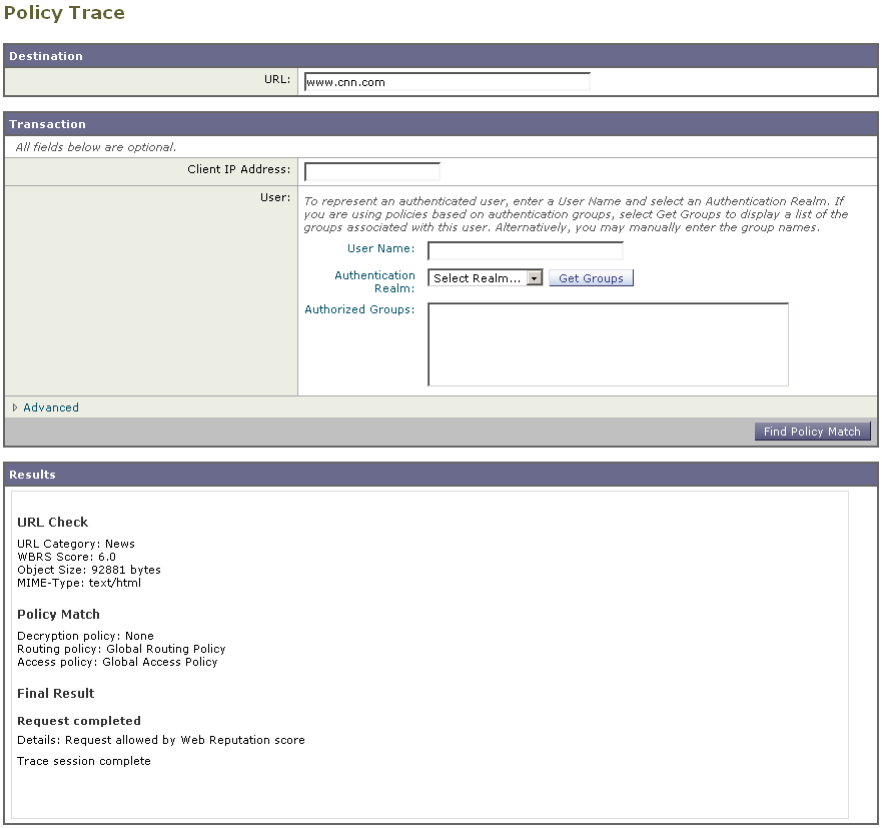

CHAPTER 6

Identities

This chapter contains the following information:

IDENTITIES OVERVIEW

To control web traffic on the network and protect your network from web based threats, AsyncOS needs to identify who is trying to access the web. Users can be identified by different criteria, such as their machine address or authenticated user name. AsyncOS can apply different actions to transactions based on the who is submitting the request.

To identify who is accessing the web, you create identities in the Web Security appliance. An identity is a policy that identifies and groups users. An identity addresses the question, "who are you?"

Identities are the only policy where you define whether or not authentication is required to access the web. However, identities do *not* specify a list of users who are *authorized* (allowed) to access the web. You specify authorized users in the other policy types.

All other policy types use an identity as the basis to determine which policy group applies to the transaction.

You might want to group the following types of users or machines:

- **A group of machine addresses in a test lab.** You can create a routing policy with this identity so requests from these machines are fetched directly from the destination server.
- **All authenticated users based on the All Realms authentication sequence.** You can create a single access policy using this identity, or you can create a different access policy for each authentication realm and configure different access control settings for users in each realm.
- **Users accessing the Web Security appliance on a particular proxy port.** You can create a routing policy using this identity that fetches content from a particular external proxy for requests that explicitly connect to the appliance on a particular proxy port.
- **All subnets trying to access a website in a user defined URL category do not require authentication.** You can create an access policy using this identity to exempt requests to particular destinations from authentication. You might want to do this for Windows update servers.

Define identities on the Web Security Manager > Identities page. For more information about creating identities, see "Creating Identities" on page 112.

Note — IronPort recommends creating fewer, more general identity groups that apply to all users or fewer, larger groups of users on the network. Then, if you need to apply more granular access control to the web, you can create multiple access and decryption policies that apply to smaller groups of authorized users.

EVALUATING IDENTITY GROUP MEMBERSHIP

When a client sends a request to a server, AsyncOS receives the request, evaluates it, and determines to which identity group it belongs.

To determine the identity group that a client request matches, AsyncOS follows a very specific process for matching the identity group membership criteria. During this process, it considers the following factors for group membership:

- **Subnet.** The client subnet must match the list of subnets in a policy group.
- **Port.** The proxy port of the request must be in the identity group's list of ports, if any are listed. For explicit forward connections, this is the port configured in the browser. For transparent connections, this is the same as the destination port.

  You might want to define identity group membership on the proxy port if you have one set of clients configured to explicitly forward requests on one port, and another set of clients configured to explicitly forward requests on a different port.

  Note — IronPort recommends only defining identity group membership by the proxy port when the appliance is deployed in explicit forward mode, or when clients explicitly forward requests to the appliance. When you define identity group membership by the proxy port when clients requests get transparently redirected to the appliance, some requests might be erroneously denied.

- **User agent.** The user agent making the request must be in the identity group's list of user agents, if any are listed. You might want to group by user agent for user agents that cannot handle authentication and you want to create an identity that does not require authentication.
- **URL category.** The URL category of the request URL must be in the identity group's list of URL categories, if any are listed. You might want to group by URL destination category if you create different authentication groups based on URL categories and want to apply them to users depending on the website categorization.
- **Authentication requirements.** If the identity group requires authentication, the client authentication credentials must match the identity group's authentication requirements. For more information about how authentication works with identity groups, see "How Authentication Affects Identity Groups" on page 106.

The information in this section gives an overview of how the appliance matches client requests to identity groups. For more details on exactly how the appliance matches client requests, see "Matching Client Requests to Identity Groups" on page 109.

AsyncOS sequentially reads through each identity group in the identity policies table. It compares the client request status to the membership criteria of the first identity group. If they match, AsyncOS assigns the identity group to the transaction.

If they do not match, AsyncOS compares the client request to the next identity group. It continues this process until it matches the client request to a user defined identity group, or if

it does not match a user defined identity group, it matches the global identity policy. When AsyncOS matches the client request to an identity group or the global identity policy, it assigns the identity group to the transaction.

If at any time during the comparison process the user fails authentication, AsyncOS terminates the request. For more information about how authentication works with identity groups, see "How Authentication Affects Identity Groups" on page 106.

After AsyncOS assigns an identity to a client request, it evaluates the request against the other policy group types. For more information, see the following locations:

- "Evaluating Access Policy Group Membership" on page 122
- "Evaluating Decryption Policy Group Membership" on page 170
- "Evaluating Routing Policy Group Membership" on page 143

How Authentication Affects Identity Groups

Requiring authentication for users can help your organization control access to the web for groups of users. AsyncOS allows you to create multiple identity groups and define the membership criteria based on authentication requirements.

When authentication is required for an identity group, a gold key icon appears next to the identity group name in the Policies table, as shown in Figure 6-1.

Figure 6-1 Identity Groups that Require Authentication

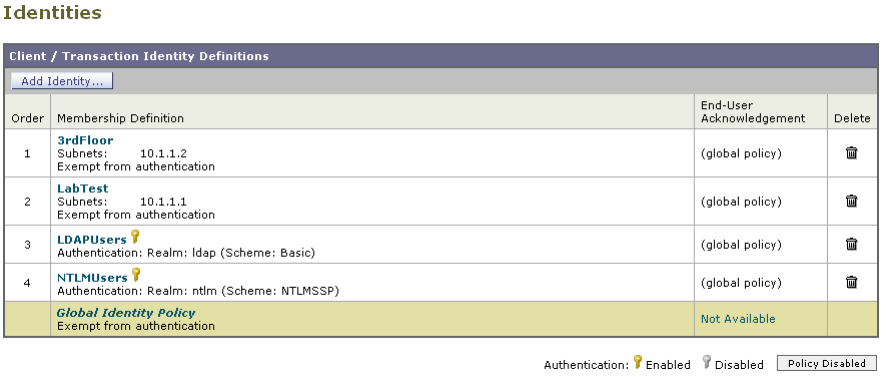

To define authentication requirements for an identity group, you can choose an authentication realm or sequence that applies to the identity group.

Note — You can specify the authorized users when you use the identity in a different type of policy group.

Consider the following rules and guidelines when creating and ordering identity groups:

- **Identity group order.** All identity groups that do not require authentication must be above identity groups that require authentication.
- **Transparent appliance deployment.** When the appliance is deployed in transparent mode with cookie-based authentication, it does not get cookie information from clients for HTTPS requests. Therefore, it cannot match an identity group requiring authentication to an HTTPS request. In this situation, any HTTPS request that does not match a user defined identity group will always match the global identity policy, *even if the global identity policy requires authentication*. This ensures that all HTTPS requests match at least one identity group.
- **Identity uniqueness.** Verify the identity group membership requirements are unique for each identity group. If two identity groups require the exact same membership, then client requests never match the lower identity group. If any access, decryption, or routing policies use the lower identity group, client requests never match those policies.
- **Global identity policy.** The global identity policy does not require authentication by default when you create an authentication realm. If you want the global identity policy to require authentication, you must assign an authentication realm, authentication sequence, or the All Realms sequence to the global identity policy.

For some examples of how AsyncOS matches client requests to an identity group for different identity policies tables, see "Example Identity Policies Tables" on page 115.

How Authentication Affects HTTPS Requests

How the appliance matches HTTPS requests with identities depends on the appliance deployment and browser configuration:

- **Explicit forward mode.** The appliance matches HTTPS requests with identity groups the same way it matches HTTP requests. For a diagram of how this occurs, see Figure 6-2 on page 110.
- **Transparent mode, but the browser is explicitly configured to use the appliance as a proxy.** The behavior is the same as when the appliance is in explicit forward mode. For a diagram of how this occurs, see Figure 6-2 on page 110.
- **Transparent mode with IP-based authentication type.** The behavior is different, depending on whether or not the HTTPS request comes from a client that has authentication information available from an earlier HTTP request:
  - **Information available from a previous HTTP request.** The behavior is the same as when the appliance is in explicit forward mode. For a diagram of how this occurs, see Figure 6-2 on page 110.
  - **No information available from a previous HTTP request.** When the appliance has no credential information for the client, then it fails the HTTPS request.
- **Transparent mode with cookie-based authentication type.** When the appliance is deployed in transparent mode with cookie-based authentication, it does not get cookie information from clients for HTTPS requests. Therefore, it cannot match an HTTPS request

to an identity group requiring authentication. In this situation, HTTPS requests can only match identity groups that do not require authentication or the global identity policy, *even if the global identity policy requires authentication*. Any HTTPS request that does not match a user defined identity group will always match the global identity policy, even if the global identity policy requires authentication. This ensures that all HTTPS requests match at least one identity group. For a diagram of how this occurs, see Figure 6-3 on page 111.

How Authentication Scheme Affects Identity Groups

You define the authentication scheme for each identity group, not at each realm or sequence. That means you can use the same NTLM realm or a sequence that contains an NTLM realm and use it in identity groups that use either the NTLMSSP, Basic, or "Basic or NTLMSSP" authentication schemes.

AsyncOS communicates which scheme(s) it supports to the client application at the beginning of a transaction. The identity group currently in use determines which scheme(s) it supports. When AsyncOS informs the client application that it supports both Basic and NTLMSSP, the client application chooses which scheme to use in the transaction.

Some client applications, such as Internet Explorer, always choose NTLMSSP when given a choice between NTLMSSP and Basic. This might cause a user to not pass authentication when all of the following conditions are true:

- The identity group uses a sequence that contains both LDAP and NTLM realms.
- The identity group uses the "Basic or NTLMSSP" authentication scheme.
- A user sends a request from an application that chooses NTLMSSP over Basic.
- The user only exists in the LDAP realm.

When this happens, AsyncOS uses the NTLMSSP scheme to authentication users in this identity group because the client requests it. However, LDAP servers do not support NTLMSSP, so no user that exists only in the specified LDAP server(s) can pass authentication in this identity group.

Therefore, when you need to use an authentication sequence that contains both LDAP and NTLM realms, consider the client applications that might try to access a URL when you configure the authentication scheme for an identity group. For example, you might want to choose Basic as the only authentication scheme for an identity group in some cases.

MATCHING CLIENT REQUESTS TO IDENTITY GROUPS

Figure 6-2 on page 110 shows how AsyncOS evaluates a client request against the identity groups under the following circumstances:

- The Web Security appliance is deployed in explicit forward mode.
- The Web Security appliance is deployed in transparent mode, but a client application is configured to explicitly forward transactions to the Web Security appliance.
- The Web Security appliance is deployed in transparent mode and it uses IP-based authentication.

Figure 6-3 on page 111 shows how AsyncOS evaluates a client request against the identity groups when the Web Security appliance is deployed in transparent mode with cookie-based authentication. The evaluation process is different for transparent, cookie-based authentication because the Web Proxy can never authenticate the user for HTTPS transactions.

Figure 6-2 Policy Group Flow Diagram for Identities - Explicit Forward and Transparent IP-Based

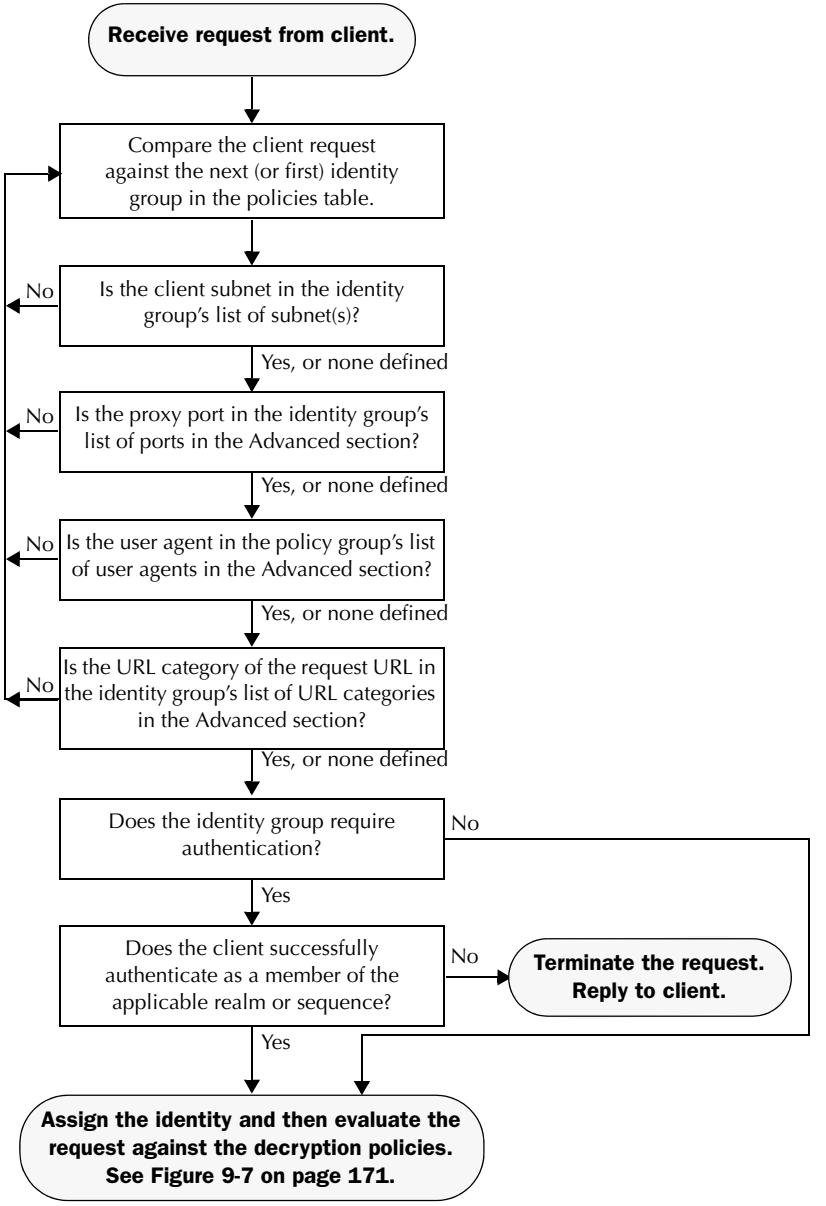

Figure 6-3 Policy Group Flow Diagram for Identities - Transparent Cookie-Based

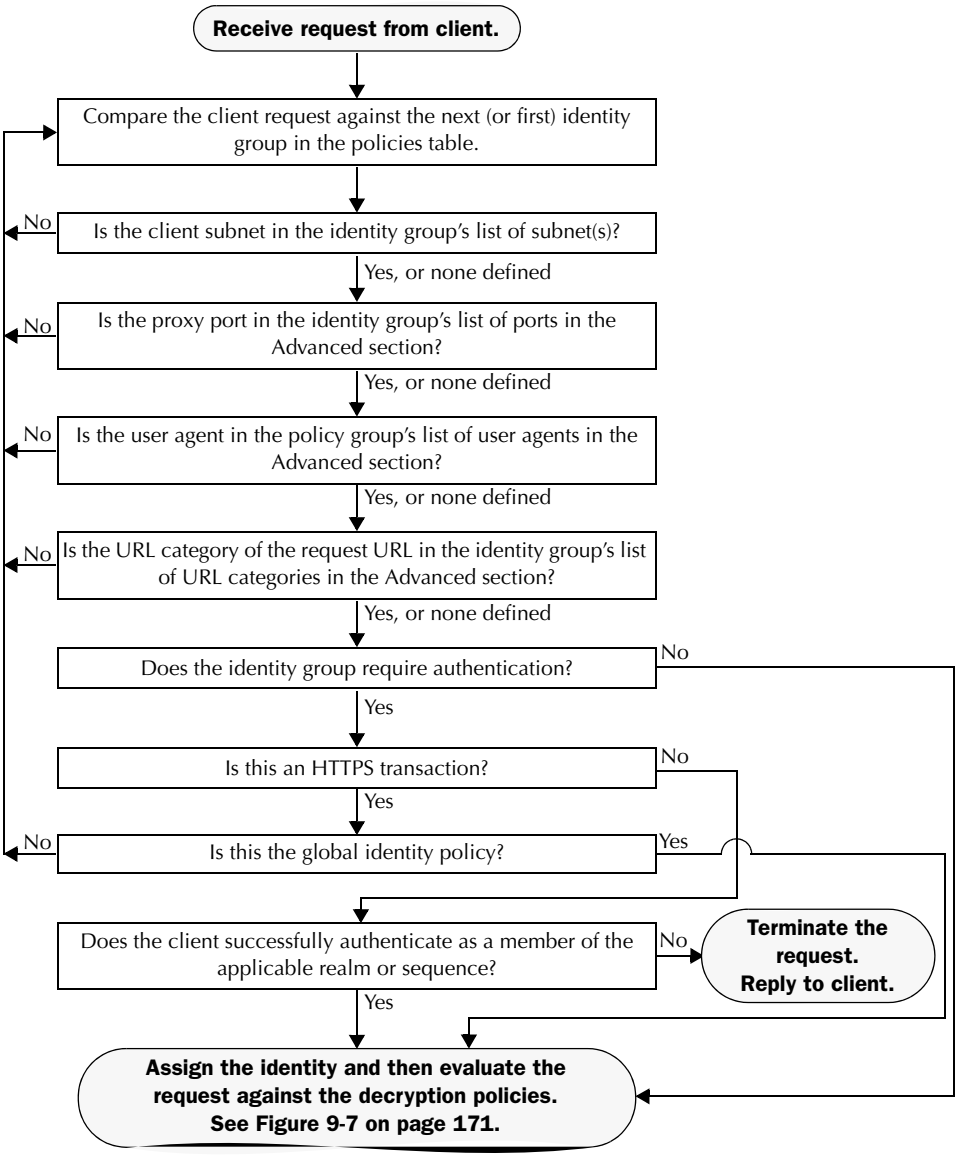

CREATING IDENTITIES

You can create identities based on combinations of several criteria, such as client subnet or the URL category of the destination site. You must define at least one criterion for identity membership. When you define multiple criteria, the client request must meet all criteria to match the identity.

For more information about how the appliance matches a client request with an identity, see "Evaluating Identity Group Membership" on page 105 and "Matching Client Requests to Identity Groups" on page 109.

You define policy group membership on the Web Security Manager > Identities page.

Note — If you delete an authentication realm or sequence, any identity that depends on the deleted realm or sequence becomes disabled.

To create an identity group:

1. Navigate to the Web Security Manager > Identities page.
2. Click **Add Identity**.

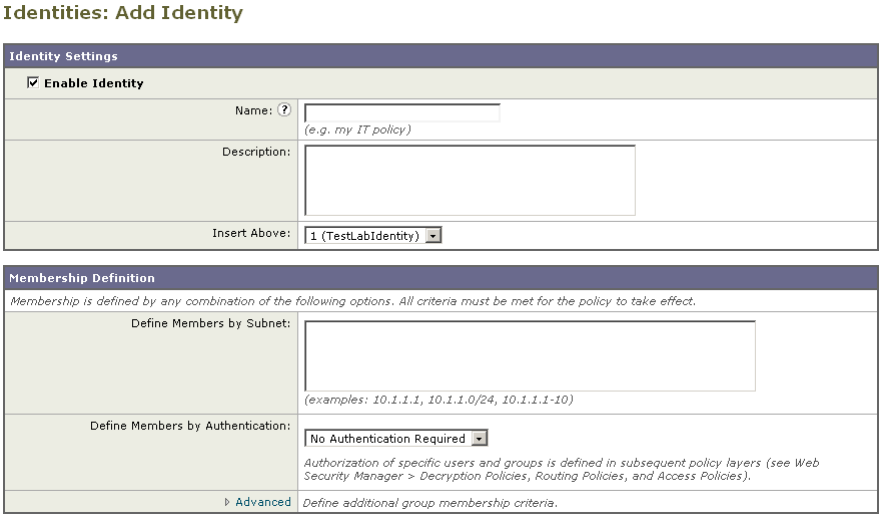

3. In the Name field, enter a name for the identity group, and in the Description field, optionally add a description.

   Note — Each identity group name must be unique.

4. In the Insert Above field, choose where in the policies table to place the identity group.

When configuring multiple identity groups, you must specify a logical order for each group. Carefully order your identity groups to ensure that correct matching occurs and position groups that do not require authentication above the first policy group that does require authentication. For more information about how authentication affects identity groups, see “How Authentication Affects Identity Groups” on page 106.

5. In the Define Members by Subnet field, enter the addresses to which this identity should apply.

   You can enter IP addresses, CIDR blocks, and subnets. Separate multiple addresses with commas.

   Note — If you do not enter an address in this field, the identity group applies to *all* IP addresses. For example, if you configure the identity to require authentication, but do not define any other settings, then the identity acts similarly to the Default Identity Policy with authentication required.

6. In the Define Members by Authentication section, choose whether or not this identity requires authentication. You can choose No Authentication Required or you can choose a defined authentication realm or sequence.
7. If you choose an NTLM authentication realm or sequence that contains an NTLM authentication realm, you can choose the authentication scheme in the Scheme field.

   Note — You can specify individual authenticated users or groups of users when you use the identity in a different type of policy group.

8. Optionally, expand the Advanced section to define additional membership requirements.

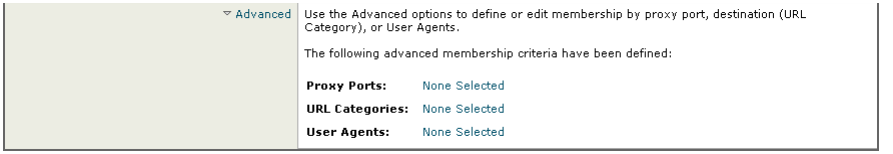

9. To define identity group membership by any of the advanced options, click the link for the advanced option and configure the option on the page that appears.

Table 6-1 describes the advanced options you can configure for identity groups.

Table 6-1 Identity Group Advanced Options

Advanced Option	Description
Proxy Ports	To define policy group membership by the proxy port used to access the Web Proxy, enter one or more port numbers in the Proxy Ports field. Separate multiple ports with commas. For explicit forward connections, this is the port configured in the browser. For transparent connections, this is the same as the destination port. You might want to define policy group membership on the proxy port if you have one set of clients configured to explicitly forward requests on one port, and another set of clients configured to explicitly forward requests on a different port. **Note:** IronPort recommends only defining policy group membership by the proxy port when the appliance is deployed in explicit forward mode, or when clients explicitly forward requests to the appliance. When you define policy group membership by the proxy port when clients requests get transparently redirected to the appliance, some requests might be denied.
URL Categories	Choose the user defined or predefined URL categories. Membership for both user defined and predefined URL categories are excluded by default, meaning the Web Proxy ignores all categories unless they are selected in the Add column.
User Agents	Choose whether or not to define policy group membership by the user agent used in the client request. You can select some commonly defined browsers, or define your own using regular expressions. Choose whether this policy group should apply to the selected user agents or to any user agent that is not in the list of selected user agents. For more information on creating user agent based policies, see "Working with User Agent Based Policies" on page 95.

10. Submit and commit your changes.

EXAMPLE IDENTITY POLICIES TABLES

This section shows some sample identity groups defined in an identity policies table and describes how AsyncOS evaluates different client requests using each identity policies table.

Example 1

Table 6-2 shows an identity policies table with three user defined identity groups. The first identity group applies to a particular subnet and does not require authentication. The second identity group applies to all subnets and requests for URLs in the "Proxies & Translators" category, and requires authentication on RealmA. The third identity group applies to all subnets, has no advanced options defined, and requires authentication on RealmA. The global identity policy applies to all subnets (by definition) and does not require authentication.

Table 6-2 Policies Table Example 1

Order	Subnet(s)	Authentication Required?	Realm or Sequence	Advanced Options
1	10.1.1.1	No	N/A	none
2	All	Yes	RealmA	URL Category is "Proxies & Translators"
3	All	Yes	RealmA	none
Global identity policy	All (by default)	No	N/A	N/A (none by default)

AsyncOS matches client requests to identity groups in this scenario differently, depending on the client's subnet and the URL category of the request:

- **Any client on subnet 10.1.1.1 for any URL.** When a client on subnet 10.1.1.1 sends a request for any URL, AsyncOS evaluates the first identity group and determines that the client subnet matches the first identity group subnet. Then it determines that no authentication is required and no advanced options are configured, so it assigns the first identity group to the transaction.
- **Any client on a subnet other than 10.1.1.1 for URLs in the "Proxies & Translators" URL category.** When a client on a subnet other than 10.1.1.1 sends a request for a URL in the "Proxies & Translators" category, AsyncOS evaluates the first identity group and determines that the client subnet is not listed in the first identity group's list of subnets. Therefore, it evaluates the second identity group, and then determines that the client subnet is listed in the second identity group's list of subnets. Then it determines that the URL in the request matches the URL category in the second identity group's advanced section. Then it determines that the second identity group requires authentication, so it

tries to authenticate the user against the authentication server(s) defined in RealmA. If the user exists in RealmA, AsyncOS assigns the second identity group to the transaction. If the user does not exist in RealmA, AsyncOS terminates the client request because the client failed authentication.

- **Any client on a subnet other than 10.1.1.1 for any URL *not* in the "Proxies & Translators" URL category.** When a client on a subnet other than 10.1.1.1 sends a request for a URL, AsyncOS evaluates the first identity group and determines that the client subnet is not listed in the first identity group's list of subnets. Therefore, it evaluates the second identity group, and then determines that the client subnet is listed in the second identity group's list of subnets. Then it determines that the URL in the request *does not* match the URL category in the second identity group's advanced section. Therefore, it evaluates the third identity group, and then determines that the client subnet is listed in the third identity group's list of subnets. The third identity group does not have any advanced options configured, so continues to compare against authentication requirements. Then it determines that the third identity group requires authentication, so it tries to authenticate the user against the authentication server(s) defined in RealmA. If the user exists in RealmA, AsyncOS assigns the third identity group to the transaction. If the user does not exist in RealmA, AsyncOS terminates the client request because the client failed authentication.

Note that in this scenario, most client requests will never match the global identity group because of the user defined identity group (the third group) that applies to all subnets, has no advanced options, and requires authentication. Any client on the network that does not match the first or second identity group will match the third identity group. The exception to this is for HTTPS requests when the appliance is in transparent mode with cookie-based authentication. Any client on a subnet other than 10.1.1.1 will match the global identity group even though it requires authentication.

Example 2

Table 6-3 shows a policies table with two user defined identity groups. The first identity group applies to all subnets, requires authentication, and specifies RealmA for authentication. The second identity group applies to all subnets, requires authentication, and specifies RealmB for authentication. Neither identity group has any advanced option configured. The global identity group applies to all subnets, requires authentication, and specifies the All Realms sequence for authentication.

Table 6-3 Policies Table Example 2

Order	Subnet(s)	Authentication Required?	Realm or Sequence	Advanced Options
1	All	Yes	RealmA	none
2	All	Yes	RealmB	none

Table 6-3 Policies Table Example 2 (Continued)

Order	Subnet(s)	Authentication Required?	Realm or Sequence	Advanced Options
Global identity policy	All	Yes	All Realms	N/A (none by default)

In this scenario, when a client sends a request for a URL, AsyncOS evaluates the first identity group and determines that the identity group applies to all subnets and has no advanced options configured. It determines that the identity group requires authentication and that the only realm specified in the identity group is RealmA. Therefore, *in order for a client on any subnet to pass authentication, it must exist in RealmA*.

When a client that exists in RealmA sends a request for a URL, the client passes authentication and AsyncOS assigns the first identity group to the transaction. When a client that does *not* exist in RealmA sends a request for a URL, the client fails authentication and AsyncOS terminates the request.

Note that when a client in RealmB sends a request for a URL, AsyncOS does *not* match the client request with the second identity group. This is because a previous identity group already applies to the same subnets (and the exact same advanced options, which in this example is none) in the second identity group and it requires authentication, but from RealmA instead. Clients in RealmB do not "fall through" to the second identity group.

If you want users in RealmB to have different access, decryption, and routing policy settings applied to them than users in RealmA, perform the following steps:

1. Create an authentication sequence that contains both RealmA and RealmB. You can choose the order of the realms in the sequence depending on your business needs.
2. Create one identity group and configure it for whichever subnets on which users in RealmA and RealmB might exist. In this example, you would configure the identity group for all subnets.
3. Configure the identity group to use the sequence you defined in step 1.
4. Create two user defined policy groups of the same type, such as access policies, and configure them both to use the identity group with the authentication sequence you defined in step 3.
5. Configure the first policy group to only apply to users in one realm, such as RealmA. You can do this by specifying a particular realm in the sequence, or by using authentication groups, or entering specific usernames.
6. Configure the second policy group to only apply to users in the other realm, such as RealmB. You can do this by specifying a particular realm in the sequence, or by using authentication groups, or entering specific usernames.

When you configure the appliance in this way, any client that sends a request for a URL must exist in either realm in the sequence (RealmA or RealmB) in order to pass authentication at the identity level. Once an identity has been assigned to the client request, AsyncOS can compare the client request against the other policy types and determine which policy group, such as an access policy group, to match and then apply those access control settings. In this example, AsyncOS matches users in RealmA with the policy group configured in step 5, and matches users in RealmB with the policy group configured in step 6.

CHAPTER 7

Access Policies

This chapter contains the following information:

ACCESS POLICIES OVERVIEW

AsyncOS for Web uses multiple web security features in conjunction with its Web Proxy and DVS engine to control web traffic, protect networks from web-based threats, and enforce organization acceptable use policies. You can define policies that determine which HTTP connections are allowed and blocked.

To configure the appliance to handle HTTP requests, perform the following tasks:

1. **Enable the Web Proxy.** To allow or block HTTP traffic, you must first enable the Web Proxy. Usually, the Web Proxy is enabled during the initial setup using the System Setup Wizard. For more information, see "Configuring the Web Proxy" on page 58.
2. **Create and configure access policy groups.** After the Web Proxy is enabled, you create and configure access policy groups to determine how to handle each request from each user. For more information, see "Access Policy Groups" on page 120.

Access Policy Groups

Access policies define how the Web Proxy handles HTTP requests and decrypted HTTPS connections for network users. You can apply different actions to specified groups of users. You can also specify which ports the Web Proxy monitors for HTTP transactions.

When the Web Proxy receives an HTTP request on a monitored port or a decrypted HTTPS connection, it compares the request to the access policy groups to determine which access policy group to apply. After it assigns the request to an access policy group, it can determine what to do with the request. For more information about evaluating policy group membership, see "Policy Group Membership" on page 90.

The Web Proxy can perform any of the following actions on an HTTP request or decrypted HTTPS connection:

- **Allow.** The Web Proxy permits the connection without interruption. Allowed connections may not have been scanned by the DVS engine.
- **Block.** The Web Proxy does not permit the connection and instead displays an end user notification page explaining the reason for the block.
- **Redirect.** The Web Proxy does not allow the connection to the originally requested destination server and instead connects to a different specified URL. You might want to redirect traffic at the appliance if your organization published the links to an internal site, but the location of the site changed since publication, or if you do not have control over the web server. For more information about redirecting traffic, see "Redirecting Traffic" on page 219.

Note — The preceding actions are final actions that the Web Proxy takes. on a client request. The Monitor action that you can configure for access policies is not a final action. For more information, see "Understanding the Monitor Action" on page 121.

After the Web Proxy assigns an access policy group to an HTTP or decrypted HTTPS request, it compares the request to the policy group's configured access control settings to determine which action to apply. You can configure multiple security components to determine how to handle HTTP and decrypted HTTPS requests for a particular policy group. For more information about the security components that you can configure and how the Web Proxy uses access policy groups to control HTTP traffic, see "Controlling Access to HTTP Traffic" on page 128.

Understanding the Monitor Action

When the Web Proxy compares a transaction to the access control settings, it evaluates the settings in order. Each access control setting can be configured to perform one of the following actions for access policies:

- Monitor
- Allow
- Block
- Redirect

All actions except Monitor are final actions that the Web Proxy applies to a transaction. A final action is an action that causes the Web Proxy to stop comparing the transaction to the rest of the access control settings.

The Monitor action is an intermediary action. The Web Proxy continues comparing the transaction to the other access control settings to determine which final action to apply.

For example, if an access policy is configured to *monitor* a suspect user agent, the Web Proxy does not make a final determination about a request from the user agent. If an access policy is configured to *block* a particular URL category, then any request to that URL category is blocked before fetching the content from the server regardless of the server's reputation score.

Note — When an access control setting matches Monitor and the transaction is ultimately allowed, the Web Proxy logs the monitored setting in the access logs. For example, when a URL matches a monitored URL category, the Web Proxy logs the URL category in the access logs.

Figure 7-3 on page 129 shows the order that the Web Proxy uses when evaluating access control settings for access policies. The flow diagram shows that the only actions applied to a transaction are the final actions: Allow, Block, and Redirect.

Note — Figure 9-9 on page 178 shows the order the Web Proxy uses when evaluating access control settings for decryption policies.

EVALUATING ACCESS POLICY GROUP MEMBERSHIP

After AsyncOS assigns an identity to a client request, AsyncOS evaluates the request against the other policy types to determine which policy group it belongs for each type. When HTTPS scanning is enabled, it applies HTTP and *decrypted* HTTPS requests against the access policies. When HTTPS scanning is not enabled, by default, it evaluates all HTTP and HTTPS requests against the access policies.

AsyncOS applies the configured policy control settings to a client request based on the client request's policy group membership.

To determine the policy group that a client request matches, AsyncOS follows a very specific process for matching the group membership criteria. During this process, it considers the following factors for group membership:

- **Identity.** Each client request either matches an identity or fails authentication and gets terminated. For more information about evaluating identity group membership, see "Evaluating Identity Group Membership" on page 105.
- **Authorized users.** If the assigned identity requires authentication, the user must be in the list of authorized users in the access policy group to match the policy group.
- **Advanced options.** You can configure several advanced options for access policy group membership. Some of the options (subnet, proxy port, and URL category) can also be defined within the identity. When you configure an advanced option at the access policy group level that is also configured within the identity, you narrow down the field of advanced options that the transaction must match.

The information in this section gives an overview of how the appliance matches client requests to access policy groups. For more details about exactly how the appliance matches client requests, see "Matching Client Requests to Access Policy Groups" on page 122.

AsyncOS sequentially reads through each policy group in the policies table. It compares the client request status to the membership criteria of the first policy group. If they match, AsyncOS applies the policy settings of that policy group.

If they do not match, AsyncOS compares the client request to the next policy group. It continues this process until it matches the client request to a user defined policy group, or if it does not match a user defined policy group, it matches the global policy group. When AsyncOS matches the client request to a policy group or the global policy group, it applies the policy settings of that policy group.

Matching Client Requests to Access Policy Groups

Figure 7-1 on page 123 shows how AsyncOS evaluates a client request against the access policy groups.

Figure 7-1 Policy Group Flow Diagram for Access Policies

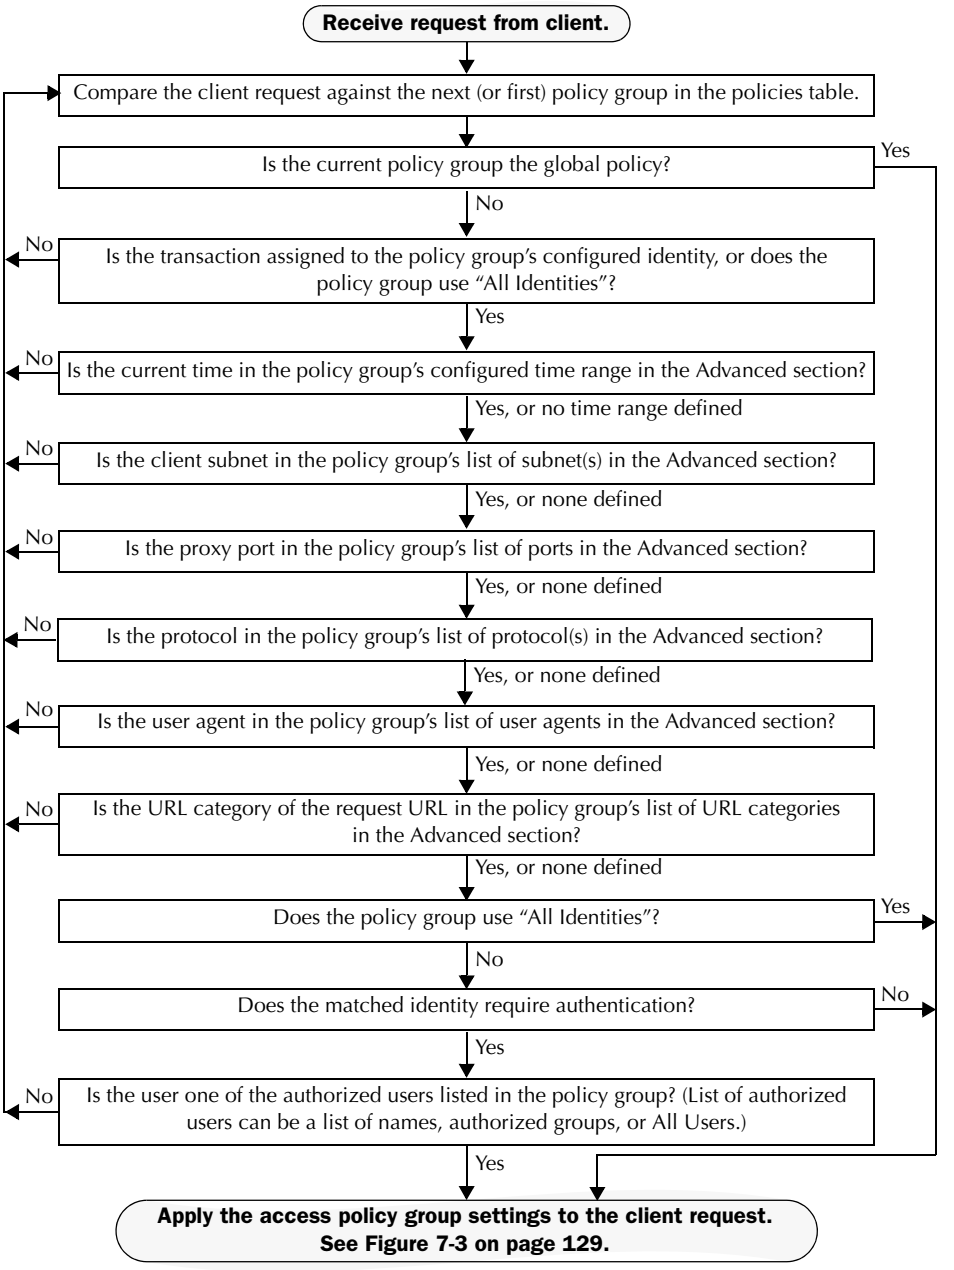

CREATING ACCESS POLICIES

You can create access policy groups based on combinations of several criteria, such as identity or the URL category of the destination site. You must define at least one criterion for policy group membership. When you define multiple criteria, the client request must meet all criteria to match the policy group.

For more information about how the appliance matches a client request with a policy group, see "Evaluating Access Policy Group Membership" on page 122 and "Matching Client Requests to Access Policy Groups" on page 122.

You define policy group membership on the Web Security Manager > Access Policies page.

To create an access policy group:

1. Navigate to the Web Security Manager > Access Policies page.
2. Click **Add Group**.
3. In the Policy Name field, enter a name for the policy group, and in the Description field, optionally add a description.
4. In the Insert Above Policy field, choose where in the policies table to place the policy group.

   When configuring multiple policy groups you must specify a logical order for each group. Carefully order your policy groups to ensure that correct matching occurs.
5. In the Policy Member Definition section, Identity Policy subsection, choose the identity group to apply to this policy group.
6. If you choose an identity that requires authentication, you can specify which users are authorized for this policy group. These users must authenticate. In the second field in the Identity Policy section, you can choose one of the following options:
   - **All users.** Choose Include all users in this realm.
   - **Specific users.** Choose Specify authorized groups and users, and enter the users in the other fields in the section. The fields that appear depend on the type of authentication used in the identity.
7. Optionally, expand the Advanced section to define additional membership requirements.

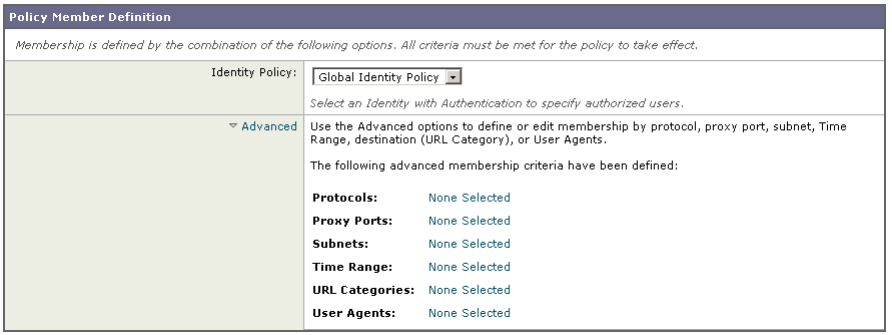

8. To define policy group membership by any of the advanced options, click the link for the advanced option and configure the option on the page that appears.

   Table 7-1 describes the advanced options you can configure for access policy groups.

Table 7-1 Access Policy Group Advanced Options

Advanced Option	Description
Protocols	Choose whether or not to define policy group membership by the protocol used in the client request. Select the protocols to include. "All others" means any protocol not listed in above this option. **Note:** When HTTPS scanning is enabled, only decryption policies apply to HTTPS transactions. You cannot define policy membership by the HTTPS protocol for access or routing policies.

Table 7-1 Access Policy Group Advanced Options (Continued)

Advanced Option	Description
Proxy Ports	Choose whether or not to define policy group membership by the proxy port used to access the Web Proxy. Enter one or more port numbers in the Proxy Ports field. Separate multiple ports with commas. For explicit forward connections, this is the port configured in the browser. For transparent connections, this is the same as the destination port. You might want to define policy group membership on the proxy port if you have one set of clients configured to explicitly forward requests on one port, and another set of clients configured to explicitly forward requests on a different port. IronPort recommends only defining policy group membership by the proxy port when the appliance is deployed in explicit forward mode, or when clients explicitly forward requests to the appliance. When you define policy group membership by the proxy port when clients requests get transparently redirected to the appliance, some requests might be denied. **Note:** The identity associated with this policy group might define identity membership by proxy port. Adding proxy ports in the policy group further narrows down the list of transactions that match this policy group.
Subnets	Choose whether or not to define policy group membership by subnet or other addresses. You can choose to use the addresses that may be defined with the associated identity, or you can enter specific addresses here. **Note:** If the associated identity defines identity membership by addresses, you must enter addresses that are a subset of the addresses defined in the identity. Adding addresses in the policy group further narrows down the list of transactions that match this policy group.
Time Range	Choose whether or not to define policy group membership by a defined time range. Choose the time range from the Time Range field and then choose whether this policy group should apply to the times inside or outside the selected time range. For more information on creating time based policies, see “Working with Time Based Policies” on page 93. For more information on creating time ranges, see “Creating Time Ranges” on page 93.

Table 7-1 Access Policy Group Advanced Options (Continued)

Advanced Option	Description
URL Categories	Choose whether or not to define policy group membership by URL categories. Select the user defined or predefined URL categories. **Note:** The identity associated with this policy group might define identity membership by URL categories. Adding URL categories in the policy group further narrows down the list of transactions that match this policy group.
User Agents	Choose whether or not to define policy group membership by the user agent used in the client request. You can select some commonly defined browsers, or define your own using regular expressions. Choose whether this policy group should apply to the selected user agents or to any user agent that is not in the list of selected user agents. For more information on creating user agent based policies, see "Working with User Agent Based Policies" on page 95.

9. Submit your changes.
10. Configure access policy group access control settings to define how AsyncOS handles transactions.

    The new access policy group automatically inherits global policy group settings until you configure options for each access control component. For more information, "Controlling Access to HTTP Traffic" on page 128.
11. Submit and commit your changes.

CONTROLLING ACCESS TO HTTP TRAFFIC

After the Web Security appliance assigns an HTTP or decrypted HTTPS request to an access policy group, it assigns the access control settings of the policy group to the request. The access control settings of the access policy group determine whether the appliance allows, blocks, or redirects the connection.

Configure access control settings for access policy groups on the Web Security Manager > Access Policies page.

Figure 7-2 shows where you can configure access control settings for the access policy groups.

Figure 7-2 Creating Secure Access Policies

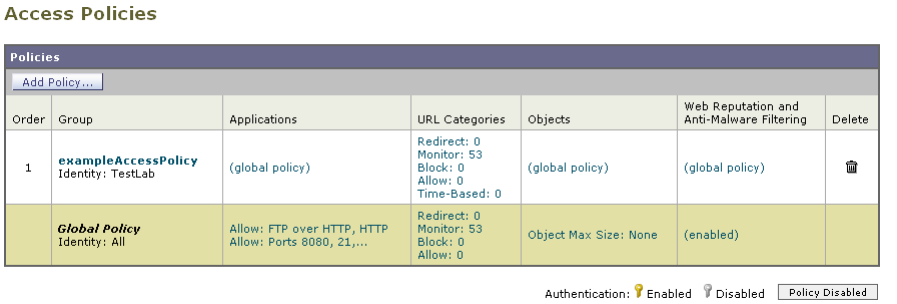

You can configure the following settings to determine what action to take on the HTTP or decrypted HTTPS request:

- **Applications.** For more information, see “Applications” on page 130.
- **URL Categories.** For more information, see “URL Categories” on page 130.
- **Objects.** For more information, see “Object Blocking” on page 131.
- **Web Reputation and Anti-Malware Filtering.** For more information, see “Web Reputation and Anti-Malware” on page 131.)

After an access policy group is assigned to a request, the access control settings for the policy group are evaluated to determine whether to allow, block, or redirect the request. For more information about assigning an access policy group to an HTTP request, see “Policy Group Membership” on page 90.

Figure 7-3 on page 129 shows how the appliance determines which action to take on an HTTP or decrypted HTTPS request after it has assigned a particular access policy to the request.

Figure 7-3 Applying Access Policy Actions

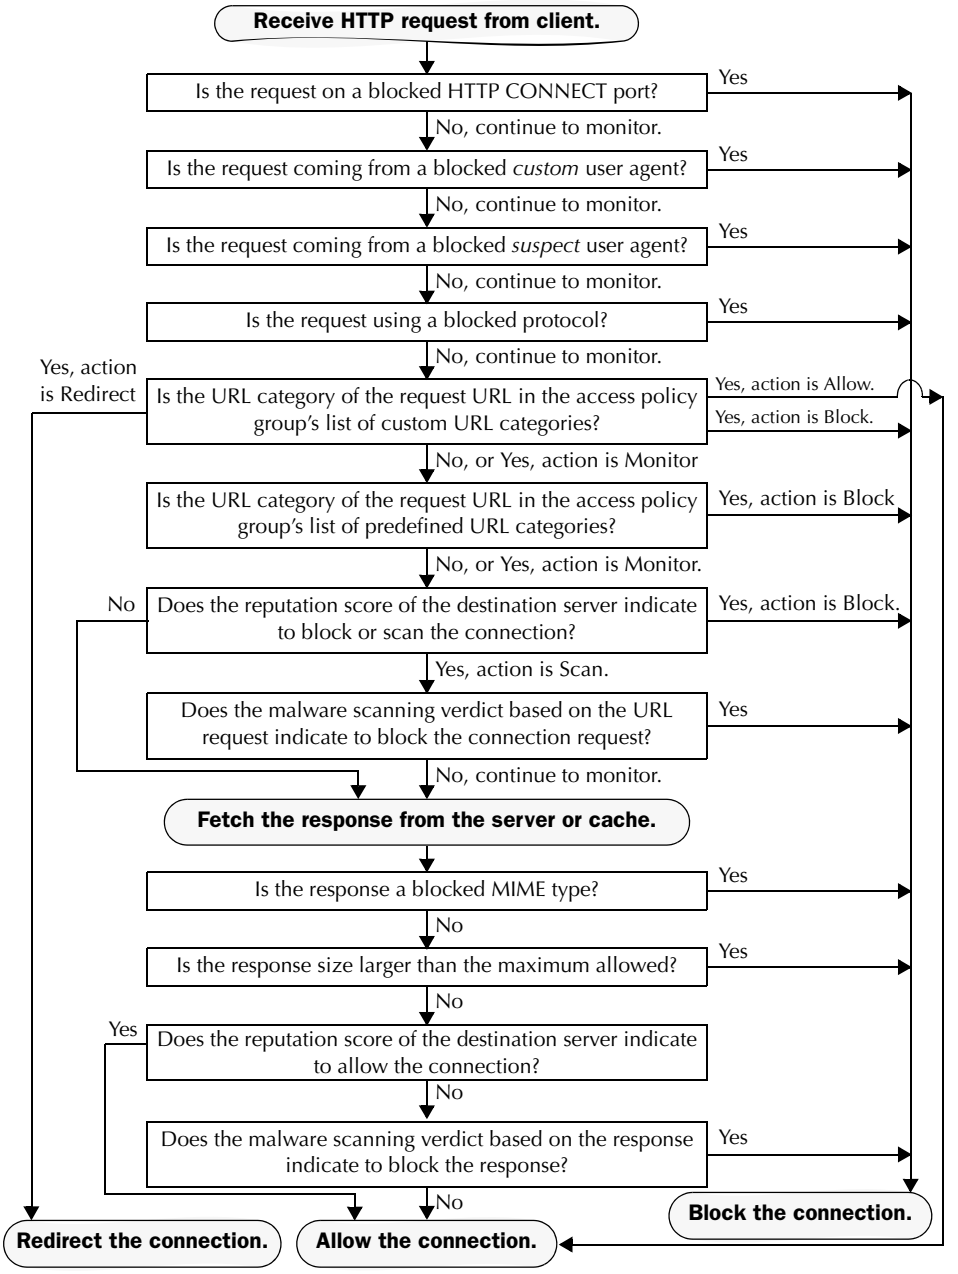

Figure 7-3 on page 129 shows two different decision points that involve the web reputation score of the destination server. The web reputation score of the server is evaluated only once, but the result is applied at two different points in the decision flow.

Applications

You can use the Applications settings on the Access Policies > Applications page to control policy group access to protocols and configure blocking for Internet applications (also known as user agents), such as instant messenging clients, web browsers, and Internet phone services. You can also configure the appliance to tunnel HTTP CONNECT requests on specific ports. With tunneling enabled, the appliance passes HTTP traffic through specified ports without evaluating it.

For more information about blocking user agents, see "Blocking Specific Applications and Protocols" on page 133.

Figure 7-4 Custom Settings for Controlling Applications

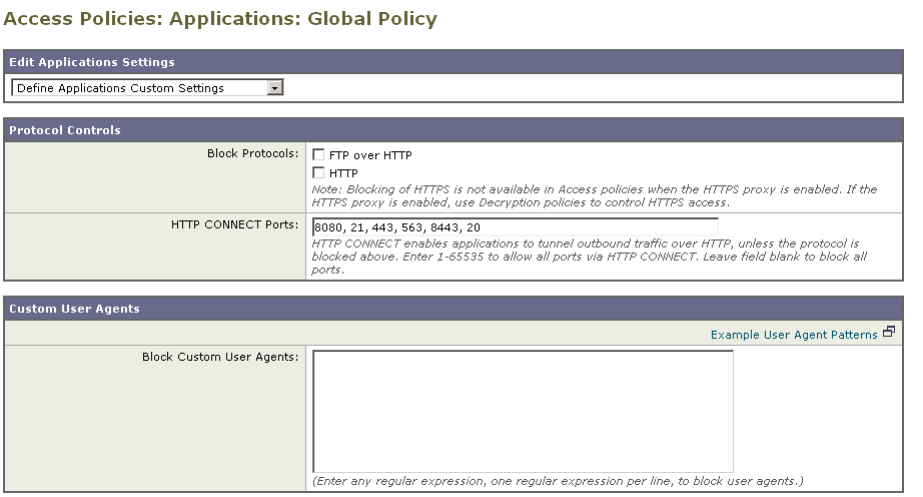

Note — When HTTPS scanning is enabled, you can only use decryption policies to control access to HTTPS transactions. You cannot configure access policies on this page to block HTTPS connections.

URL Categories

IronPort URL Filters allow you to configure how the appliance handles a transaction based on the URL category of a particular HTTP request. Using a predefined category list, you can choose to monitor or block content by category. You can also create custom URL categories

and choose to allow, monitor, block, or redirect traffic for a website in the custom category. You can use custom URL categories to create block and allow lists based on destination.

For more information about enabling URL filters and working with URL categories, see "URL Filters" on page 207.

Object Blocking

You can use the settings on the Access Policies > Objects page to configure the proxy to block file downloads based on file characteristics, such as file size and file type. For more information about blocking a specific object or MIME-type, see "Blocking Specific Applications and Protocols" on page 133.

Figure 7-5 Blocking Object Types

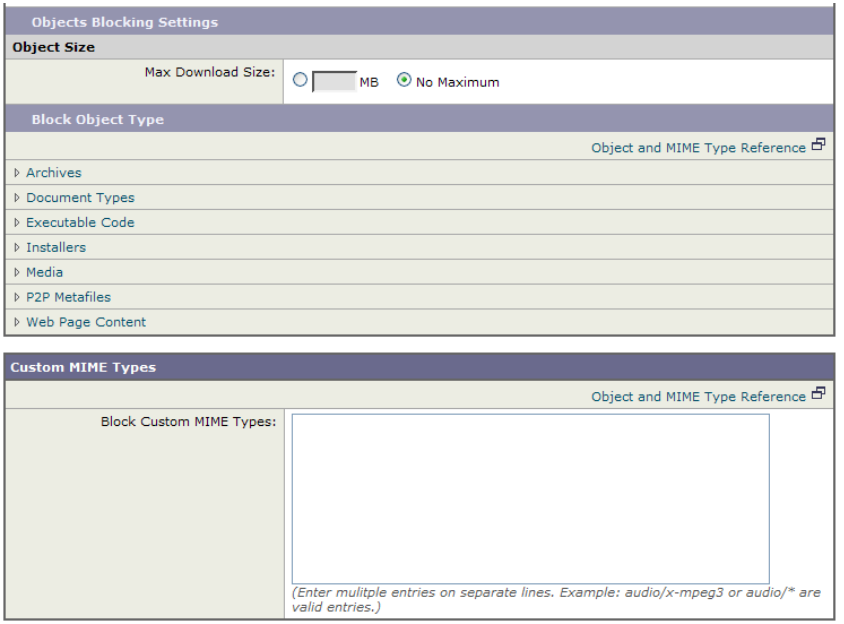

Web Reputation and Anti-Malware

The Web Reputation and Anti-Malware Filtering policy inherits global settings respective to each component. To customize filtering and scanning for a particular policy group, you can use the Web Reputation and Anti-Malware Settings pull-down menu to customize monitoring or blocking for malware categories based on malware scanning verdicts and to customize web reputation score thresholds.

For more information about configuring web reputation scores, see "Configuring Web Reputation Scores" on page 233.

For more information about configuring anti-malware settings, see “Configuring Anti-Malware Scanning” on page 246.

BLOCKING SPECIFIC APPLICATIONS AND PROTOCOLS

AOL Messenger, BitTorrent, Skype—the Web Security appliance can control and block access to these types of applications. You can configure how the appliance manages these kinds of applications based on the port being used:

- **Port 80.** You can control how the Web Security appliance manages these applications using access policies, but only as they are accessed via HTTP tunneling on port 80.
- **Ports other than 80.** You can block these applications on other ports by using the L4 Traffic Monitor.

Use the Web Security Manager > Access Policies page to manage access and monitoring for these types of applications on a more granular (per policy) level. Use the L4 Traffic Monitor to manage access and monitoring on a more global basis.

Blocking on Port 80

To block access to these types of applications where port 80 is used, you can use the Web Security Manager > Access Policies page. The Access Policies page provides several methods for blocking access. You can block access by clicking on any of the following columns for a particular policy group:

- Applications
- URL Categories
- Objects

You can block access to predefined URL categories such as Chat and Peer-to-Peer, or create your own custom URL categories. You can block specific applications based on their "agent patterns" or signatures.

You can apply some or all of these methods on various access policies by creating additional access policy groups. For details on how to create additional access policy groups, see "Creating Access Policies" on page 124.

Policy: Applications

You can create a rule that blocks a particular user agent based on its pattern using Regular Expressions.

You block access to applications based on their agent pattern similarly for the different access policies:

- **User defined policies** — On the Web Security Manager > Access Policies page, click the value in the Applications column for the desired policy. Choose Define Applications Custom Settings.
- **Global Policy** — On the Web Security Manager > Access Policies page, click the value in the Applications column for the Global Policy.

Once you view the Access Policies: Applications: *Policy_Name* page, add user agent patterns (also called signatures) to the Block Custom User Agents section of the page.

Figure 7-6 Entering Agent Patterns to Block

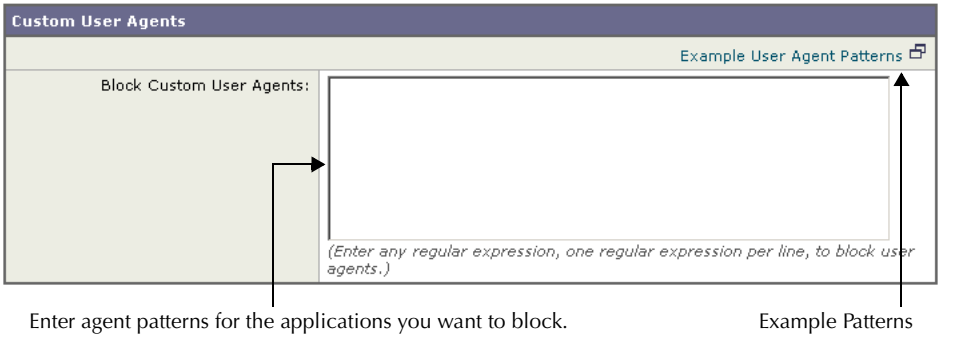

Note — You can click the Example User Agent Patterns link for a list of some example user agent patterns.

Table 7-2 provides a list of common patterns.

Table 7-2Common Application Agent Patterns

Application	Search in Setting	HTTP header	Signature
AOL Messenger	Request headers	User-Agent	Gecko/
BearShare	Response header	Server	Bearshare
BitTorrent	Request headers	User-Agent	BitTorrent
eDonkey	Request headers	User-Agent	e2dk
Gnutella	Request headers	User-Agent	Gnutella Gnucleus
Kazaa	Request headers	P2P-Agent	Kazaa Kazaaclient:
Kazaa	Request headers	User-Agent	KazaClient Kazaaclient:
Kazaa	Request headers	X-Kazaa-Network	KaZaA
Morpheus	Response header	Server	Morpheus
MSN Messenger	Request headers	User-Agent	MSN Messenger

Table 7-2Common Application Agent Patterns (Continued)

Application	Search in Setting	HTTP header	Signature
Trillian	Request headers	User-Agent	Trillian/
Windows Messenger	Request headers	User-Agent	MSMSGS
Yahoo Messenger	Request headers	Host	msg.yahoo.com
Yahoo Messenger	Request headers	User-Agent	ymsgr

This is not a comprehensive list, as signatures change occasionally, and new applications are developed. You can find additional signatures at various websites, including the following websites:

- http://www.user-agents.org/
- http://www.useragentstring.com/pages/useragentstring.php
- http://www.infosyssec.com/infosyssec/security/useragentstrings.shtml

Note — IronPort Systems does not maintain, verify, or support the user agent listings at any of these websites.

Policy: URL Categories

You can specify categories of URLs to block, including the predefined "Chat" and "Peer-to-Peer" categories. You can also add specific custom URL categories should you want to add a URL that is not already included in the predefined categories. You may then add the custom category to the list of blocked URLs.

For more information about using URL Categories, see "URL Categories" on page 130.

Policy: Objects

You can block some Peer-to-Peer files directly, via the Access Policies: Objects: Global Policy page.

On the Web Security Manager > Access Policies page, click on the value in the Objects column for the desired policy.

In the Block Object Type section, check any boxes in the P2P Metafiles group. You can add custom MIME (Multipurpose Internet Mail Extensions) types by entering them in the Custom MIME Types field. For example, entering the `application/x-zip` signature blocks ZIP archive files.

Blocking on Ports Other Than 80

If these applications are using ports other than 80, you may want to block access to a specific server or block of IP addresses to which the client must connect. To manage these applications on other ports, use the L4 Traffic Monitor. The L4 Traffic monitor allows you to

restrict access on specific ports. However, the restriction is global, so it will apply to all traffic on that port.

CHAPTER 8

Working with External Proxies

This chapter contains the following topics:

WORKING WITH EXTERNAL PROXIES OVERVIEW

The Web Security appliance is a proxy-compatible device, and is easily deployed within an existing proxy environment. However, you must place the appliance downstream from existing proxy servers, meaning closer to the clients. Positioning the appliance between internal clients and an existing proxy server enables the appliance to detect and stop malware on client machines before it reaches the existing proxy server, and protects malware from the Internet reaching your internal network.

You can configure the appliance to work with multiple existing, upstream proxies. Use the Network > Upstream Proxies page to define upstream proxies or to modify existing settings. You define groups of proxies, and you can configure the appliance to use load balancing and failover features when connecting to multiple proxies.

After defining proxy groups, you can create routing policies to determine whether the Web Proxy connects to the server identified by the client or to a member of one the proxy groups.

For more information about using routing policies to route transactions, see “Routing Traffic to Upstream Proxies” on page 139. For more information about defining external proxies, see “Adding External Proxy Information” on page 141.

ROUTING TRAFFIC TO UPSTREAM PROXIES

When the Web Proxy does not deliver a response from the cache, it can direct client requests directly to the destination server or to an external proxy on the network. You use routing policies to create rules that indicate when and to where to direct transactions. A routing policy determines to where to pass the client request, either to another proxy (as defined by the proxy group) or to the destination server. It addresses the question, "from where to fetch content?" You might want to create routing policies if you have a highly distributed network.

Figure 8-1 shows routing policies on the Web Security Manager > Routing Policies page.

Figure 8-1 Routing Policies

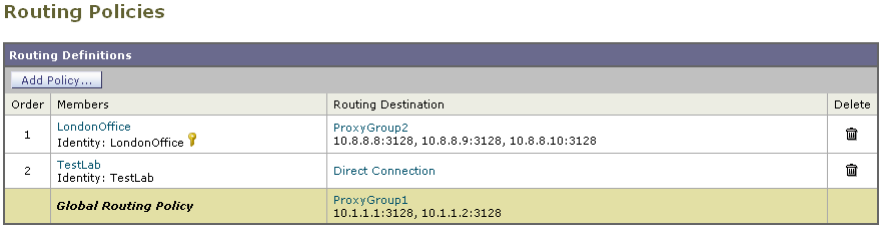

When you define multiple external proxies in a proxy group, the Web Proxy can use load balancing techniques to distribute requests to different proxies defined in the group. You can choose the following load balancing techniques:

- **None (failover).** The Web Proxy directs transactions to one external proxy in the group. It tries to connect to the proxies in the order they are listed. If one proxy cannot be reached, the Web Proxy attempts to connect to the next one in the list.
- **Fewest connections.** The Web Proxy keeps track of how many active requests are with the different proxies in the group and it directs a transaction to the proxy currently servicing the fewest number of connections.
- **Hash based.** The Web Proxy uses a hash function to distribute requests to the proxies in the group. The hash function uses the proxy ID and URL as inputs so that requests for the same URL are always directed to the same external proxy.
- **Least recently used.** The Web Proxy directs a transaction to the proxy that least recently received a transaction if all proxies are currently active. This setting is similar to round robin except the Web Proxy also takes into account transactions a proxy has received by being a member in a different proxy group. That is, if a proxy is listed in multiple proxy groups, the "least recently used" option is less likely to overburden that proxy.
- **Round robin.** The Web Proxy cycles transactions equally among all proxies in the group in the listed order.

For information about creating routing policies, see "Creating Routing Policies" on page 145.

Note — If your network contains an upstream proxy that does not support FTP connections, then you must create a routing policy that applies to all identities and to just FTP requests. Configure that routing policy to directly connect to FTP servers or to connect to a proxy group whose proxies all support FTP connections.

ADDING EXTERNAL PROXY INFORMATION

To define external proxy information, you create a proxy group. A proxy group is an object that defines a list of proxies and their connection information and the load balancing technique to use when distributing requests to proxies in the group. You can create multiple proxy groups and can define multiple proxies within a group.

AsyncOS for Web allows you to enter the same proxy server information multiple times into the same proxy group. You might want to include the same proxy server multiple times to allow unequal load distribution among the proxies in the proxy group.

Note — You can only specify one existing proxy during the System Setup Wizard. AsyncOS creates a proxy group with one proxy using the information you enter in the System Setup Wizard. You can specify additional proxies in the web interface after initial setup.

To create a proxy group:

1. Navigate to Network > Upstream Proxies, and click **Add Group**.

   The Add Upstream Proxy Group page appears.

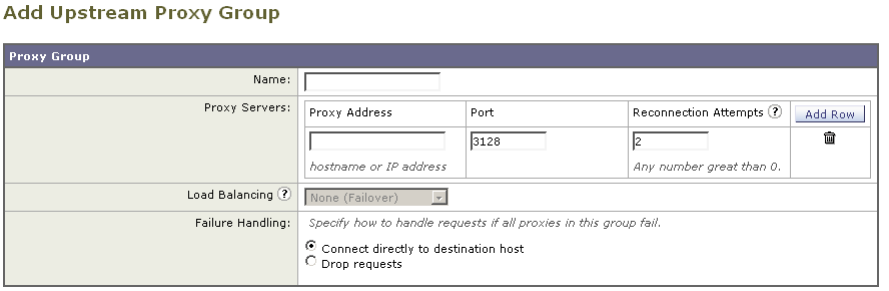

2. Enter a name for the proxy group in the Name field.
3. In the Proxy Servers section, define at least one external proxy.
   a. In the Proxy Address field, enter the host name or IP address of the proxy server.
   b. In the Port field, enter the port number used to access the proxy.
   c. In the Reconnection Attempts field, enter the number of times the Web Proxy should try to connect to the proxy server before ignoring it.
   d. Optionally, you can define another proxy server by clicking Add Row.
4. In the Load Balancing field, choose the method the Web Proxy should use to distribute transactions to the proxies when the group contains multiple proxies.

   For more information about the load balancing options, see “Routing Traffic to Upstream Proxies” on page 139.

5. In the Failure Handling field, choose how the Web Proxy should handle transactions when all proxies in the group fail.
6. Submit and commit your changes.

EVALUATING ROUTING POLICY GROUP MEMBERSHIP

After AsyncOS assigns an identity to a client request, AsyncOS evaluates the request against the other policy types to determine which policy group it belongs for each type. Any request that does not get terminated due to failed authentication gets evaluated against the routing policies to determine from where to fetch the data.

Once AsyncOS assigns a routing policy group to a request, it fetches the content from the location configure for the policy group, either from a configured proxy group or directly from the server.

To determine the policy group that a client request matches, AsyncOS follows a very specific process for matching the group membership criteria. During this process, it considers the following factors for group membership:

- **Identity.** Each client request either matches an identity or fails authentication and gets terminated. For more information about evaluating identity group membership, see "Evaluating Identity Group Membership" on page 105.
- **Authorized users.** If the assigned identity requires authentication, the user must be in the list of authorized users in the routing policy group to match the policy group.
- **Advanced options.** You can configure several advanced options for routing policy group membership. Some of the options (subnet, proxy port, and URL category) can also be defined within the identity. When you configure an advanced option at the routing policy group level that is also configured within the identity, you narrow down the field of advanced options that the transaction must match.

The information in this section gives an overview of how the appliance matches client requests to routing policy groups. For more details about exactly how the appliance matches client requests, see "Matching Client Requests to Routing Policy Groups" on page 143.

AsyncOS sequentially reads through each policy group in the policies table. It compares the client request status to the membership criteria of the first policy group. If they match, AsyncOS applies the policy settings of that policy group.

If they do not match, AsyncOS compares the client request to the next policy group. It continues this process until it matches the client request to a user defined policy group, or if it does not match a user defined policy group, it matches the global policy group. When AsyncOS matches the client request to a policy group or the global policy group, it applies the policy settings of that policy group.

Matching Client Requests to Routing Policy Groups

Figure 8-2 on page 144 shows how AsyncOS evaluates a client request against the routing policy groups.

Figure 8-2 Policy Group Flow Diagram for Routing Policies

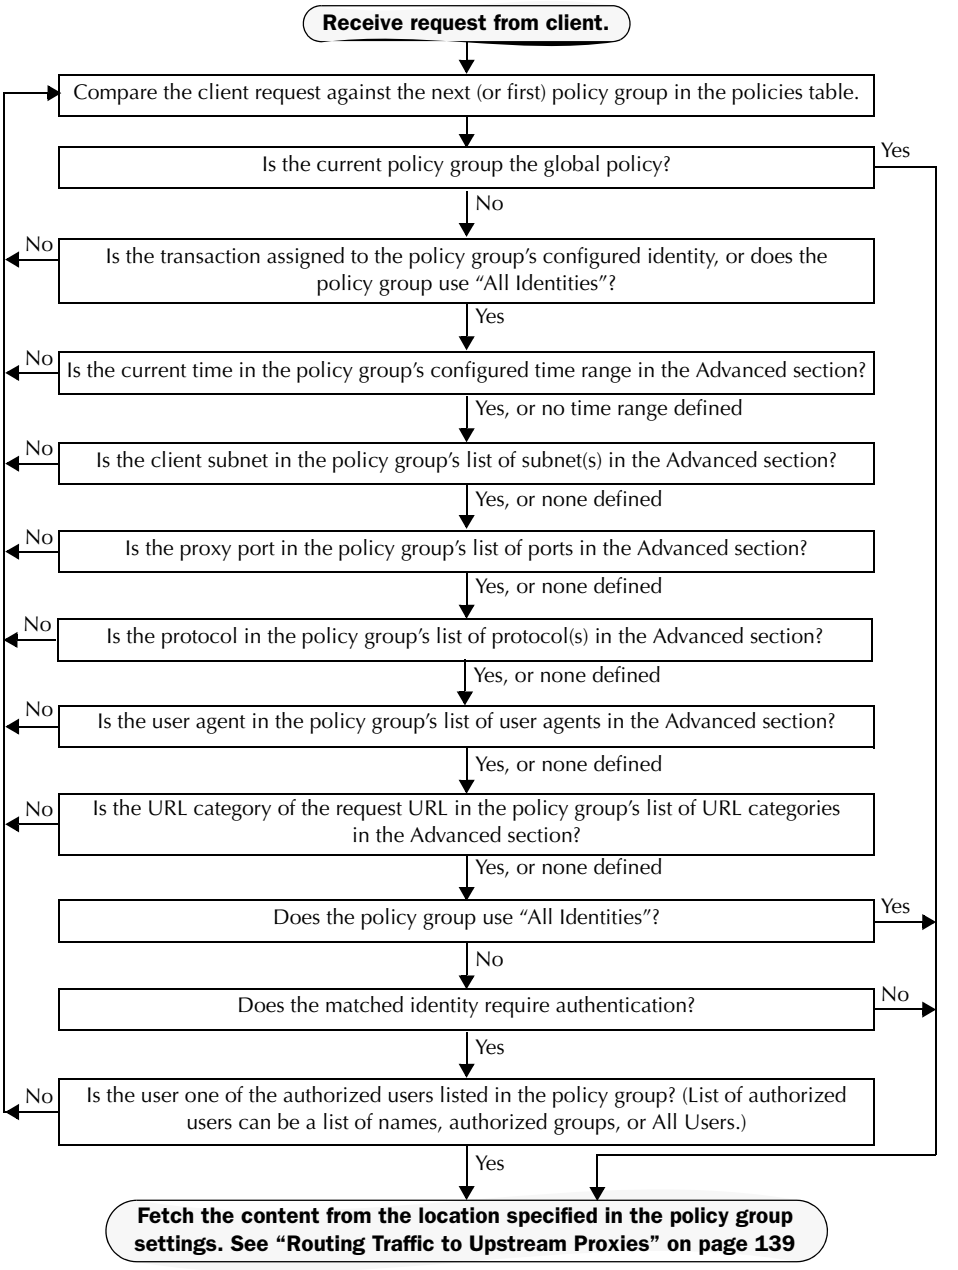

CREATING ROUTING POLICIES

You can create routing policy groups based on combinations of several criteria, such as identity or the port used to access the Web Proxy. You must define at least one criterion for policy group membership. When you define multiple criteria, the client request must meet all criteria to match the policy group.

For more information about how the appliance matches a client request with a policy group, see "Evaluating Routing Policy Group Membership" on page 143 and "Matching Client Requests to Routing Policy Groups" on page 143.

You define policy group membership on the Web Security Manager > Routing Policies page.

To create a routing policy group:

1. Navigate to the Web Security Manager > Routing Policies page.
2. Click **Add Group**.
3. In the Policy Name field, enter a name for the policy group, and in the Description field, optionally add a description.
4. In the Insert Above Policy field, choose where in the policies table to place the policy group.

   When configuring multiple policy groups you must specify a logical order for each group. Carefully order your policy groups to ensure that correct matching occurs.
5. In the Policy Member Definition section, Identity Policy subsection, choose the identity group to apply to this policy group.
6. If you choose an identity that requires authentication, you can specify which users are authorized for this policy group. These users must authenticate. In the second field in the Identity Policy section, you can choose one of the following options:
   - **All users.** Choose Include all users in this realm.
   - **Specific users.** Choose Specify authorized groups and users, and enter the users in the other fields in the section. The fields that appear depend on the type of authentication used in the identity.
7. Optionally, expand the Advanced section to define additional membership requirements.

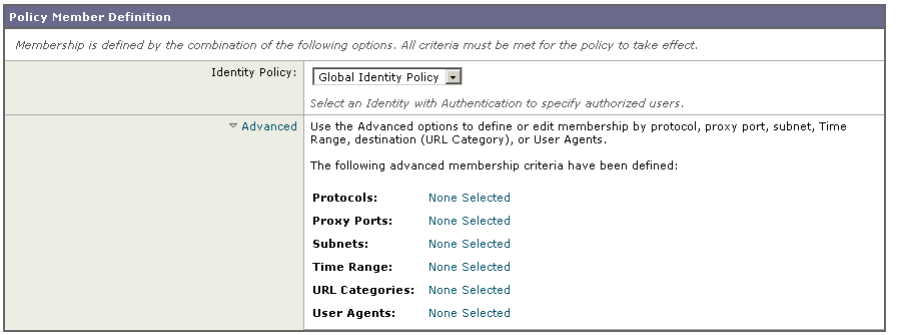

8. To define policy group membership by any of the advanced options, click the link for the advanced option and configure the option on the page that appears.

   Table 8-1 describes the advanced options you can configure for policy groups.

Table 8-1 Policy Group Advanced Options

Advanced Option	Description
Protocols	Choose whether or not to define policy group membership by the protocol used in the client request. Select the protocols to include. "All others" means any protocol not listed in above this option. **Note:** When HTTPS scanning is enabled, only decryption policies apply to HTTPS transactions. You cannot define policy membership by the HTTPS protocol for access or routing policies.

Table 8-1 Policy Group Advanced Options (Continued)

Advanced Option	Description
Proxy Ports	Choose whether or not to define policy group membership by the proxy port used to access the Web Proxy. Enter one or more port numbers in the Proxy Ports field. Separate multiple ports with commas. For explicit forward connections, this is the port configured in the browser. For transparent connections, this is the same as the destination port. You might want to define policy group membership on the proxy port if you have one set of clients configured to explicitly forward requests on one port, and another set of clients configured to explicitly forward requests on a different port. IronPort recommends only defining policy group membership by the proxy port when the appliance is deployed in explicit forward mode, or when clients explicitly forward requests to the appliance. When you define policy group membership by the proxy port when clients requests get transparently redirected to the appliance, some requests might be denied. **Note:** The identity associated with this policy group might define identity membership by proxy port. Adding proxy ports in the policy group further narrows down the list of transactions that match this policy group.
Subnets	Choose whether or not to define policy group membership by subnet or other addresses. You can choose to use the addresses that may be defined with the associated identity, or you can enter specific addresses here. **Note:** If the associated identity defines identity membership by addresses, you must enter addresses that are a subset of the addresses defined in the identity. Adding addresses in the policy group further narrows down the list of transactions that match this policy group.
Time Range	Choose whether or not to define policy group membership by a defined time range. Choose the time range from the Time Range field and then choose whether this policy group should apply to the times inside or outside the selected time range. For more information on creating time based policies, see “Working with Time Based Policies” on page 93. For more information on creating time ranges, see “Creating Time Ranges” on page 93.

Table 8-1 Policy Group Advanced Options (Continued)

Advanced Option	Description
URL Categories	Choose whether or not to define policy group membership by URL categories. Select the user defined or predefined URL categories. **Note:** The identity associated with this policy group might define identity membership by URL categories. Adding URL categories in the policy group further narrows down the list of transactions that match this policy group.
User Agents	Choose whether or not to define policy group membership by the user agent used in the client request. You can select some commonly defined browsers, or define your own using regular expressions. Choose whether this policy group should apply to the selected user agents or to any user agent that is not in the list of selected user agents. For more information on creating user agent based policies, see "Working with User Agent Based Policies" on page 95.

9. Submit your changes.
10. Configure routing policy group access control settings to define how AsyncOS handles transactions.

    The new policy group automatically inherits global policy group settings until you configure options for each access control component. For more information, see "Routing Traffic to Upstream Proxies" on page 139.
11. Submit and commit your changes.

CHAPTER 9

Decryption Policies

This chapter contains the following information:

DECRYPTION POLICIES OVERVIEW

HTTPS is a web protocol that acts as a secure form of HTTP. HTTPS encrypts HTTP requests and responses before they are sent across the network. Common thinking is that any connection to a site using HTTPS is "safe." HTTPS connections are secure, not safe, and they do not discriminate against malicious or compromised servers. HTTPS is a secure way to complete legitimate transactions, but more dangerously, it is a secure way to download malware which can infect your network.

Not being able to inspect HTTPS traffic makes the network vulnerable to the following risks:

- **Secure site hosting malware.** Spammers and phishers can create legitimate looking websites that are only reachable through an HTTPS connection. Some users may mistakenly trust the web server because it requires an HTTPS connection, resulting in intentional and unintentional downloaded malware.
- **Malware from HTTPS web applications.** Some malware can infect the network from legitimate web applications, such as secure email clients, by downloading attachments.
- **Secure anonymizing proxy.** Some web servers offer a proxy service over an HTTPS connection that allows users to circumvent acceptable use policies. When users on the network use a secure proxy server outside the network, they can access any website, regardless of its web reputation or malware content.

The appliance uses both IronPort URL Filters and IronPort Web Reputation Filters to make intelligent decisions about when to decrypt HTTPS connections. With this combination, administrators and end users are not forced to make a trade-off between privacy and security.

You can define HTTPS policies that determine if an HTTPS connection can proceed without examination or whether the appliance should act as an intermediary, decrypting the data passing each way and applying access policies to the data as if it were a plaintext HTTP transaction.

To configure the appliance to handle HTTPS requests, you must perform the following tasks:

1. **Enable HTTPS scanning.** To monitor and decrypt HTTPS traffic, you must first enable HTTPS scanning. For more information, see "Enabling HTTPS Scanning" on page 166.
2. **Create and configure decryption policy groups.** Once HTTPS scanning is enabled, you can create and configure decryption policy groups to determine how to handle each request from each user. For more information, see "Decryption Policy Groups" on page 151.
3. **Import custom root certificates (optional).** Optionally, you can import one or more custom root certificates so the Web Proxy can recognize additional trusted root certificate authorities used by HTTPS servers. For more information, see "Importing a Trusted Root Certificate" on page 180.

This book uses many terms from digital cryptography. This book also includes sections with background information about HTTPS and digital cryptography for reference only. For a list of

the terms and definitions used in this book, see "Digital Cryptography Terms" on page 153. For an overview of HTTPS the protocol, see "HTTPS Basics" on page 155.

Decryption Policy Groups

Decryption policies define how the appliance should handle HTTPS connection requests for users on the network. You can apply different actions to specified groups of users. You can also specify which ports the appliance should monitor for HTTPS transactions.

When a client makes an HTTPS request on a monitored secure port, the appliance compares the request to the decryption policy groups to determine in which decryption policy group the request belongs. Once it assigns the request to a decryption policy group, it can determine what to do with the connection request. For more information about evaluating policy group membership, see "Policy Group Membership" on page 90.

The appliance can perform any of the following actions on an HTTPS connection request:

- **Drop.** The appliance drops the connection and does not pass the connection request to the server. The appliance does not notify the user that it dropped the connection. You might want to drop connections to third party proxies that allow users on the network bypass the organization's acceptable use policies.
- **Pass through.** The appliance passes through the connection between the client and the server without inspecting the traffic content. You might want to pass through connections to trusted secure sites, such as well known banking and financial institutions.
- **Decrypt.** The appliance allows the connection, but inspects the traffic content. It decrypts the traffic and applies access policies to the decrypted traffic as if it were a plaintext HTTP connection. By decrypting the connection and applying access policies, you can scan the traffic for malware. You might want to decrypt connections to third party email providers, such as gmail or hotmail. For more information about how the appliance decrypts HTTPS traffic, see "Decrypting HTTPS Traffic" on page 160.

Note — The actions above are final actions the Web Proxy takes on an HTTPS request. The "Monitor" action you can configure for decryption policies is not a final action. For more information, see "Understanding the Monitor Action" on page 152.

Once the appliance assigns a decryption policy to an HTTPS connection request, it evaluates the request against the policy group's configured access control settings to determine which action to take. You can configure URL filter and web reputation settings to determine how to handle HTTPS requests for a particular policy group. For more information about how the appliance uses decryption policy groups to control HTTPS traffic, see "Controlling HTTPS Traffic" on page 176.

Note — IronPort recommends creating fewer, more general decryption policy groups that apply to all users or fewer, larger groups of users on the network. Then, if you need to apply more granular control to decrypted HTTPS traffic, use more specific access policy groups. For more information about access policy groups, see "Access Policies" on page 119.

For information about creating and using policy groups, see “Working with Policies” on page 83.

Note — The next two sections contain information about digital cryptography and HTTPS for reference only.

Personally Identifiable Information Disclosure

If you choose to decrypt an end-user’s HTTPS session, then the Web Security appliance access logs and reports may contain personally identifiable information. IronPort recommends that Web Security appliance administrators take care when handling this sensitive information.

You also have the option to configure how much URI text is stored in the logs using the `advancedproxyconfig` CLI command and the `HTTPS` subcommand. You can log the entire URI, or a partial form of the URI with the query portion removed. However, even when you choose to strip the query from the URI, personally identifiable information may still remain.

Understanding the Monitor Action

When the Web Proxy evaluates the access control settings against a transaction, it evaluates the settings in a particular order. Each access control setting can be configured to one of the following actions for decryption policies:

- Monitor
- Drop
- Pass through
- Decrypt

All actions except Monitor are final actions the Web Proxy applies to a transaction. A final action is an action that causes the Web Proxy to stop evaluating the transaction against other access control settings.

Monitor is an intermediary action that indicates the Web Proxy should continue evaluating the transaction against the other access control settings to determine which final action to ultimately apply.

For example, if a decryption policy is configured to monitor invalid server certificates, the Web Proxy makes no final decision on how to handle the HTTPS transaction if the server has an invalid certificate. If a decryption policy is configured to block servers with a low web reputation score, then any request to a server with a low reputation score is dropped without considering the URL category actions.

Figure 9-9 on page 178 shows the order the Web Proxy uses when evaluating access control settings for decryption policies. Looking at the flow diagram, you can see that the only actions applied to a transaction are the final actions listed above: Drop, Pass Through, and Decrypt.

Note — Figure 7-3 on page 129 shows the order the Web Proxy uses when evaluating access control settings for access policies.

DIGITAL CRYPTOGRAPHY TERMS

To understand how encryption and decryption works, you need to understand a little bit about cryptographic encoding techniques. Figure 9-1 describes some terms used in cryptography that are discussed in this chapter.

Table 9-1 Cryptography Terms and Definitions

Term	Definition
Certificate authority	An entity which issues digital certificates for use by other parties. Certificate authorities are sometimes referred to as trusted third parties. Certificate authorities are typically commercial companies that charge for their services. However, some institutions and governments have their own certificate authorities, and some offer their services for free.
Cipher	An algorithm used for encoding and decoding text to make it unreadable to any system without the appropriate key. Ciphers work with keys to encode or decode text.
Ciphertext	Encoded text after a cipher has been applied to it.
Digital certificate	An electronic document that identifies and describes an organization that has been verified and signed by a trusted organization called a certificate authority. A digital certificate is similar in concept to an "identification card." SSL uses certificates to authenticate servers. For more information about digital certificates, see "Digital Certificates" on page 157.
Digital signature	A checksum that verifies that a message was created by the stated author and was not altered since its creation.
Key	A numeric parameter used by a cipher to encode or decode text.
Plaintext or cleartext	Message text in its original form, before it gets encoded by a cipher.
Public key cryptography	A system that uses two different keys for encoding and decoding text where one key is publicly known and available and the other key is private. With public key cryptography, anyone can send an encoded message to a server that has publicized its public key, but only the recipient server can decode the message with its private key. This is also known as asymmetric key cryptography.

Table 9-1 Cryptography Terms and Definitions (Continued)

Term	Definition
Public key infrastructure (PKI)	An arrangement that binds public keys with respective user identities by means of a certificate authority. X.509 is a standard that is an example PKI. X.509 specifies standards for public key certificates and an algorithm for validating certification paths.
Private key cryptography	A system that uses the same key for encoding and decoding text. Because both sides of the transaction need the same key, they need a secure way to communicate which key to use in a particular communication session. Usually, they set up secure communication using public key cryptography and then generate a temporary symmetric key to use for the rest of the session. This is also known as symmetric key cryptography.
Root certificate	A certificate that is the topmost certificate in a certificate tree structure. All certificates below the root certificate inherit the trustworthiness of the root certificate. Root certificates can be unsigned public key certificates or self-signed certificates.
Self-signed certificate	A digital certificate where the certificate authority is the same as the certificate creator.

HTTPS BASICS

HTTPS is a web protocol that acts as a secure form of HTTP. HTTPS is secure because the HTTP request and response data is encrypted before it is sent across the network. HTTPS works similarly to HTTP, except that the HTTP layer is sent on top of a security layer using either Secure Sockets Layer (SSL) or Transport Layer Security (TLS). SSL and TLS are very similar, so this User Guide uses "SSL" to refer to both SSL and TLS, unless otherwise specified.

Figure 9-1 shows the different OSI network layers for HTTPS and HTTP. It shows that HTTPS is the HTTP protocol at the application layer over SSL or TLS at the security layer.

Figure 9-1 HTTPS and HTTP OSI Layers

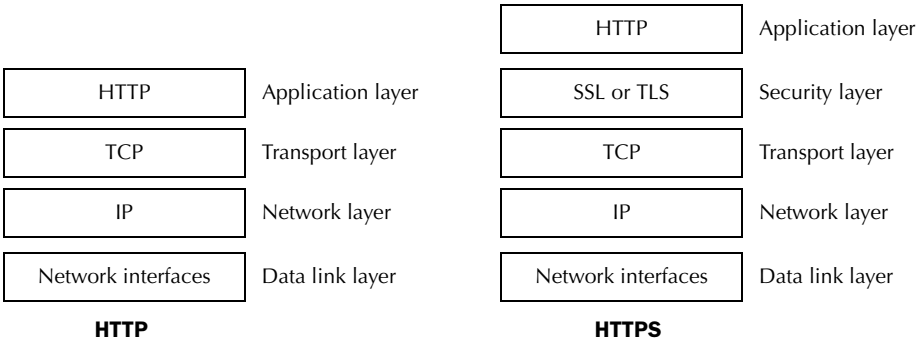

The URL typically determines whether the client application should use HTTP or HTTPS to contact a server:

- **http://*servername*.** The client application opens a connection to the server on port 80 by default and sends HTTP commands in plaintext.
- **https://*servername*.** The client application opens a connection to the server on port 443 by default and starts to engage in the SSL "handshake" to establish a secure connection between the client and server. Once the secure connection is established, the client application sends encrypted HTTP commands. For more information about the SSL handshake, see "SSL Handshake" on page 155.

SSL Handshake

The SSL "handshake" is a set of steps a client and server engage in using the SSL protocol to establish a secure connection between them. The client and server must complete the following steps before they can send and receive encrypted HTTP messages:

1. **Exchange protocol version numbers.** Both sides must verify they can communicate with compatible versions of SSL or TLS.
2. **Choose a cipher that each side knows.** First, the client advertises which ciphers it supports and requests the server to send its certificate. Then, the server chooses the strongest cipher from the list and sends the client the chosen cipher and its digital certificate.

3. **Authenticate the identity of each side.** Typically, only the server gets authenticated while the client remains unauthenticated. The client validates the server certificate. For more information about certificates and using them to authenticate servers, see "Digital Certificates" on page 157.
4. **Generate temporary symmetric keys to encrypt the channel for this session.** The client generates a session key (usually a random number), encrypts it with the server's public key, and sends it to the server. The server decrypts the session key with its private key. Both sides compute a common master secret key that will be used for all future encryption and decryption until the connection closes.

DIGITAL CERTIFICATES

A digital certificate is an electronic document that identifies and describes an organization, and that has been verified and signed by a trusted organization. A digital certificate is similar in concept to an identification card, such as a driver's license or a passport. The trusted organization that signs the certificate is also known as a certificate authority.

Certificates allow a client to know that it is talking to the organization it thinks it is talking to. When a server certificate is signed by a well-known or trusted authority, the client can better assess how much it trusts the server.

X.509 is a standard example of a public key infrastructure (PKI). X.509 specifies standards for certificates and an algorithm for validating certification paths. The Web Security appliance uses the X.509 standard.

X.509 certificates contain the following information:

- Subject's identity, such as the name of a person, server, or organization
- Certificate validity period
- Certificate authority who is vouching for the certificate
- Digital signature of the certificate created by the certificate authority using its private key
- Public key of the subject

For an example digital certificate you can view from a web browser, see "Working with Root Certificates" on page 162.

Although anyone can create a digital certificate, not everyone can get a well-respected certificate authority to vouch for the certificate's information and sign the certificate with its private key. For more information about validating the certificate authority in a digital certificate, see "Validating Certificate Authorities" on page 157.

Validating Certificate Authorities

The X.509 standard allows certificate authorities to issue digital certificates that are signed by other certificate authorities. Due to this system, there is a hierarchy of certificate authorities in a tree structure.

The top-most certificate authorities in the tree structure are called root certificates. Root certificates are not signed by a separate certificate authority because they are at the top of the tree structure. Therefore, by definition, all root certificates are self-signed certificates. The certificate authority listed in the root certificate is the certificate creator.

All certificates below the root certificate inherit the trustworthiness of the root certificate. For example, if CertificateAuthorityABC is a trusted certificate authority and it signs the certificate for certificate authority CertificateAuthorityXYZ, then CertificateAuthorityXYZ is automatically a trusted certificate authority.

Figure 9-2 shows the certification path for a certificate viewed in a web browser.

Figure 9-2 Certification Path Example

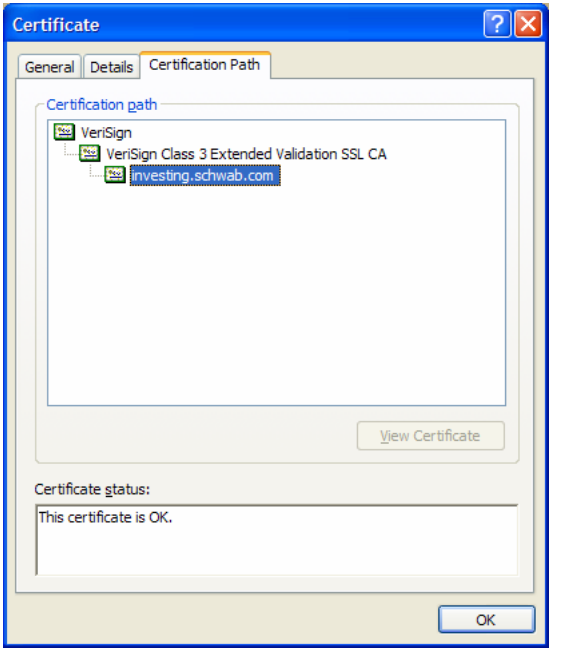

In Figure 9-2, the certificate for the URL investing.schwab.com was signed by certificate authority "VeriSign Class 3 Extended Validation SSL CA," which in turn was signed by certificate authority VeriSign.

By definition, root certificates are always trusted by applications that follow the X.509 standard. The Web Security appliance uses the X.509 standard.

Standard web browsers ship with a set of trusted root certificates. The list of root certificates is updated regularly. You can view the root certificates installed on the web browser.

For example, to view the root certificates installed with Mozilla Firefox 2.0, go to Tools > Options > Advanced > Encryption > View Certificates. To view the root certificates installed with Internet Explorer 7, go to Tools > Internet Options > Content > Certificates > Trusted Root Certification Authorities.

In Figure 9-2, the VeriSign certificate is a root certificate that shipped with the web browser.

The Web Security appliance also installs with a set of trusted root certificates. However, you can upload additional root certificates that the Web Proxy deems to be trusted. For more information about this, see "Importing a Trusted Root Certificate" on page 180.

Validating Digital Certificates

Certificates can be valid or invalid. A certificate may be in invalid for different reasons. For example, the current time may be before or after the certificate validity period, the root authority in the certificate may not be recognized, or the organization name may not match the public key.

The Web Security appliance verifies that a server certificate is valid before it inspects and decrypts an HTTPS connection from a server. You can configure how the appliance handles connections to servers with invalid certificates. The appliance can perform one of the following actions for invalid server certificates:

- **Drop.** The appliance drops the connection and does not notify the client. This is the most restrictive option.
- **Decrypt.** The appliance allows the connection, but inspects the traffic content. It decrypts the traffic and applies access policies to the decrypted traffic as if it were a plaintext HTTP connection. For more information about how the appliance decrypts HTTPS traffic, see "Decrypting HTTPS Traffic" on page 160.
- **Monitor.** The appliance does not drop the connection, and instead it continues comparing the server request with the decryption policy groups. This is the least restrictive option.

  Note — When an invalid server certificate is monitored, the errors in the certificate are maintained and passed along to the end-user.

For more information about configuring the appliance to handle invalid server certificates, see "Enabling HTTPS Scanning" on page 166.

DECRYPTING HTTPS TRAFFIC

The request and response data is encrypted for HTTPS connections before it is sent across the network. Because the data is encrypted, third parties can view the data, but cannot decrypt it to read its contents without the private key of the HTTPS server.

Figure 9-3 shows an HTTPS connection between a client and a HTTPS server.

Figure 9-3 HTTPS Connection

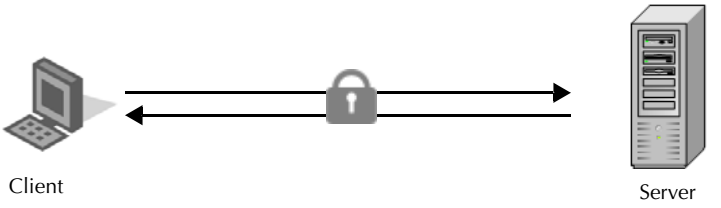

The Web Security appliance does not have access to the server's private key, so in order to inspect the traffic between the client and the server, it must intercept the connection and break the connection into two separate connections. The appliance acts as an intermediary between the client and the server pretending to be the server to the client, and the client to the server. This is sometimes referred to as being the "man in the middle."

Figure 9-4 shows an HTTPS connection between a client and a HTTPS server that goes through the Web Security appliance.

Figure 9-4 HTTPS Connection Decrypted by the Web Security Appliance

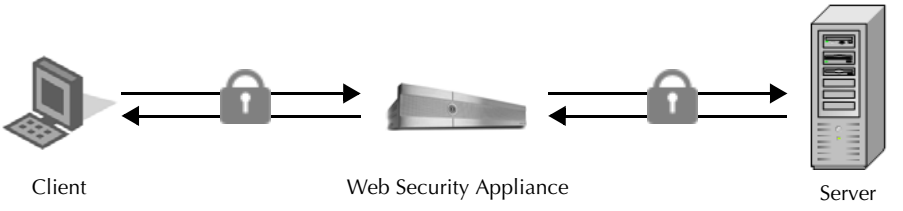

Notice that in Figure 9-4, there are two different HTTPS connections, one between the client and the appliance, and one between the appliance and the server. The appliance performs the SSL handshake twice, once with the client and again with the server:

- **SSL handshake with the server.** When the appliance performs the SSL handshake with the server, it acts as if it were the client sending a request to the server. After it establishes a secure connection with the server, it can begin receiving the encrypted data. Because it acts as the client and participates in the SSL handshake, it has agreed upon a temporary symmetric key with the server so it can decrypt and read the data the server sends. Also, the appliance receives the server's digital certificate.
- **SSL handshake with the client.** When the appliance performs the SSL handshake with the client, it acts as if it were the requested server providing data the client requests. In order

to perform the SSL handshake with the client, it must send the client its own digital certificate. However, the client expects the certificate of the requested server, so the appliance mimics the requested server's certificate by specifying a root certificate authority uploaded or configured by an appliance administrator.

For more information about how the server mimics the server's certificate, see "Mimicking the Server Digital Certificate" on page 161.

Note — Because the appliance signs the server certificate with a different root certificate authority and sends that to the client, you must verify the client applications on the network recognize the root certificate authority. For more information, see "Working with Root Certificates" on page 162.

After the two separate HTTPS connections are established, the following actions occur:

1. Encrypted data is received from the server.
2. The temporary, symmetric key negotiated with the server is used to decrypt the data.
3. Access policies are applied to the decrypted traffic as if it were a plaintext HTTP connection. For more information about access policies, see "Access Policies" on page 119.
4. Assuming the access policy group allows the client to receive the data, the data is encrypted using the temporary, symmetric key negotiated with the client.
5. Encrypted data is sent to the client.

Note — No decrypted data is cached. However, access logs for decrypted HTTP transactions are saved to disk.

Mimicking the Server Digital Certificate

When the appliance performs the SSL handshake with the client, it mimics the server digital certificate and sends the new certificate to the client. To mimic the server digital certificate, it reuses most field values and changes some field values.

The mimicked certificate is the same as the server certificate except for the following fields:

- **Issuer.** The issuer comes from the generated or uploaded root certificate configured in the appliance.
- **Signature Algorithm.** This field is always "sha1WithRSAEncryption" or "dsaWithSHA1" depending upon on whether the root certificate the appliance uses contains an RSA or DSA key.
- **Public Key.** The appliance replaces the public key in the original certificate with a public key it generates that matches bit strength from the original certificate and for which it has a matching private key generated as well. For example, if the server certificate uses a 2048 bit RSA key, the appliance generates a new 2048 bit RSA key.
- **X509v3 Extensions.** All X509v3 extensions are removed *except* for the following:

- Basic Constraints
- Subject Alternative Name
- Key Usage
- Subject Key Identifier
- Extended Key Usage

For example, the appliance removes the Authority Key Identifier and the Authority Information Access X509v3 extensions.

Working with Root Certificates

The Web Security appliance mimics the HTTPS server to which a client originally sent a connection request. In order to establish a secure connection with the client pretending to be the requested server, the appliance must send a server certificate to the client signed by a root certificate authority configured in the appliance.

When you enable HTTPS scanning on the appliance, you can configure the root certificate information that the appliance uses to sign its server certificates. You can enter root certificate information in the following ways:

- **Generate.** You can enter some basic organization information and then click a button so the appliance generates the rest of the certificate and a matching key. You might want to generate a certificate and key when your organization does not have a certificate and key in use, or when it wants to create a new and unique certificate and key.
- **Upload.** You can upload a certificate file and its matching key file created outside of the appliance. You might want to upload a certificate and matching key file if the clients on the network already have the root certificates on their machines.
  The certificate and key files you upload must be in PEM format. DER format is not supported. For more information about convert a DER formatted certificate or key to PEM format, see "Converting Certificate and Key Formats" on page 164.

  Note — The certificate you upload must contain "basicConstraints=CA:TRUE" to work with Mozilla Firefox browsers. This constraint allows Firefox to recognize the root certificate as a trusted root authority.

For more information about how to generate or upload a certificate and key, see "Enabling HTTPS Scanning" on page 166.

However, typically, the root certificate information you generate or upload in the appliance is not listed as a trusted root certificate authority in client applications. By default, when users send HTTPS requests, they will see a warning message from the client application informing them that there is a problem with the website's security certificate. Usually, the error message says that the website's security certificate was not issued by a trusted certificate authority or the website was certified by an unknown authority.

Note — You can also upload an intermediate certificate that has been signed by a root certificate authority. When the Web Proxy mimics the server certificate, it sends the uploaded

certificate along with the mimicked certificate to the client application. That way, as long as the intermediate certificate is signed by a root certificate authority that the client application trusts, the application will trust the mimicked server certificate, too. You might want to upload an intermediate certificate if your organization uses its own root certificate authority, but does not want to upload the root certificate to the Web Security appliance for security reasons.

Figure 9-5 on page 163 shows an example error message when a users sends an HTTPS request through Netscape Navigator.

Figure 9-5 Unknown Certificate Authority Error Message

Typically, users can view the certificate and use the information in the certificate to choose whether or not to allow the secure connection with this website. In Figure 9-5, you can view the certificate contents by clicking **Examine Certificate**.

Figure 9-6 on page 164 shows an example root certificate issued by the appliance.

Figure 9-6 Certificate Issued by Web Security Appliance

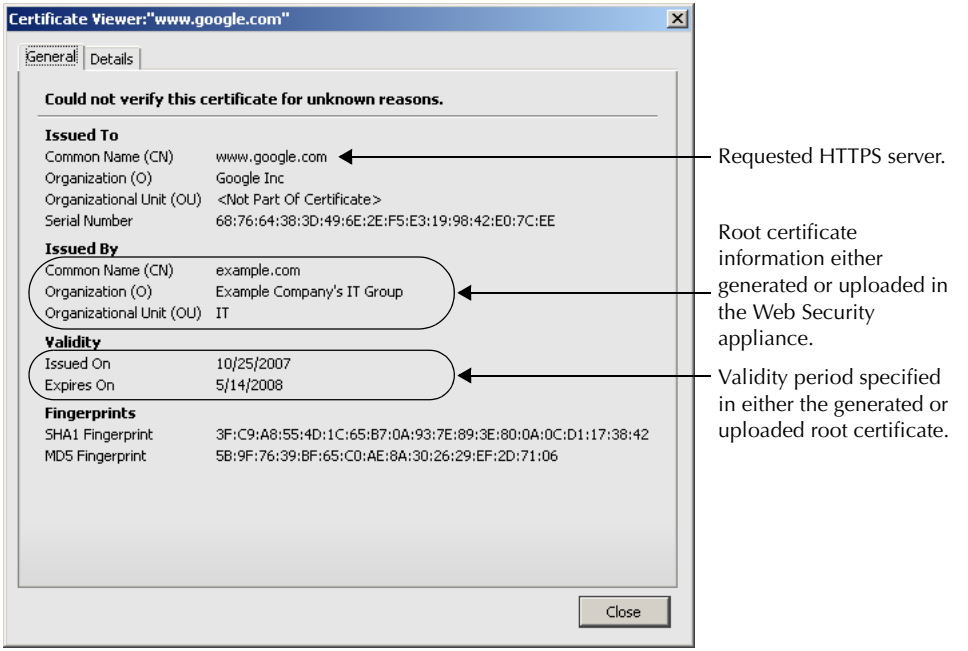

You can choose how to handle the root certificates issued by the Web Security appliance:

- **Inform users to accept the root certificate.** You can inform the users in your organization what the new policies are at the company and tell them to accept the root certificate supplied by the organization as a trusted source.
- **Add the root certificate to client machines.** You can add the root certificate to all client machines on the network as a trusted root certificate authority. This way, the client applications automatically accept transactions with the root certificate. To verify you distribute the root certificate the appliance is using, you can download the root certificate from the Security Services > HTTPS Proxy page. Click **Edit Settings**, and then click the Download Certificate link for either the generated or uploaded certificate.

  You might want to download the root certificate from the appliance if a different person uploaded the root certificate to the appliance and you want to verify you distribute the same root certificate to the client machines.

Converting Certificate and Key Formats

The root certificate file and its matching key file you upload to the appliance must be in PEM format. DER format is not supported. However, you can convert certificates and keys in DER format into the PEM format before uploading them. For example, you can use OpenSSL to convert the format.

Use the following OpenSSL command to convert a DER formatted certificate file to a PEM formatted certificate file:

```
openssl x509 -inform DER -in cert_in_DER -outform PEM -out
out_file_name
```

You can also convert key files in DER format into the PEM format by running a similar OpenSSL command.

For RSA keys, use the following command:

```
openssl rsa -inform DER -in key_in_DER -outform PEM -out out_file_name
```

For DSA keys, use the following command:

```
openssl dsa -inform DER -in key_in_DER -outform PEM -out out_file_name
```

For more information about using OpenSSL, see the OpenSSL documentation, or visit http://openssl.org.

ENABLING HTTPS SCANNING

To monitor and decrypt HTTPS traffic, you must enable HTTPS scanning on the Security Services > HTTPS Proxy page. When you enable HTTPS scanning, you must configure what the appliance uses for a root certificate when it sends self-signed server certificates to the client applications on the network. You can upload a root certificate and key that your organization already has, or you can configure the appliance to generate a certificate and key with information you enter.

Once HTTPS scanning is enabled, all HTTPS policy decisions are handled by decryption policies. You can no longer define access and routing policy group membership by HTTPS, nor can you configure access policies to block HTTPS transactions. If some access and routing policy group memberships are defined by HTTPS and if some access policies block HTTPS, then when you enable HTTPS scanning those access and routing policy groups become disabled. You can choose to enable the policies at any time, but all HTTPS related configurations are removed.

Note — When you upload a certificate to the Web Security appliance, verify it is a signing certificate and not a server certificate. A server certificate cannot be used as a signing certificate, so decryption does not work when you upload a server certificate.

For more information about root certificates, see “Working with Root Certificates” on page 162.

Also on this page, you can configure what the appliance does with HTTPS traffic when the server certificate is invalid.

Note — For information on importing a custom root authority certificate, see “Importing a Trusted Root Certificate” on page 180.

To enable HTTPS scanning:

1. Navigate to the Security Services > HTTPS Proxy page, and click **Enable and Edit Settings**.

   The HTTPS Proxy License Agreement appears.

2. Read the terms of the HTTPS Proxy License Agreement, and click **Accept**.

   The Edit HTTPS Proxy Settings page appears.

Edit HTTPS Proxy Settings

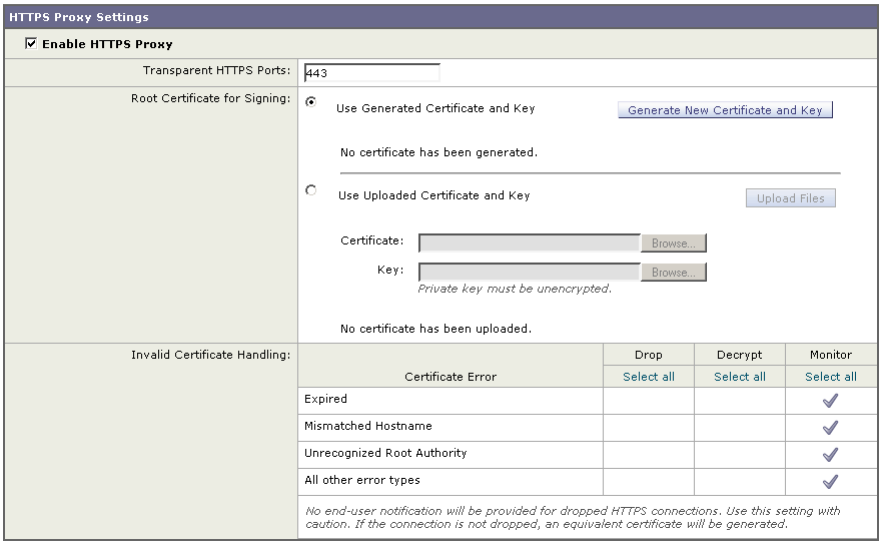

3. Verify the Enable HTTPS Proxy field is enabled.
4. Enter the ports the appliance should check for HTTPS traffic in the Transparent HTTPS Ports field. Port 443 is entered by default.

   Note — This field only appears when the appliance is deployed in transparent mode.

5. Choose which root certificate to use for signing self-signed certificates the appliance sends to clients:

   - **Generated certificate and key.** Go to step 6 on page 167.
   - **Uploaded certificate and key.** Go to step 7 on page 168.

   For more information about how the appliance uses these root certificates, see “Working with Root Certificates” on page 162.

   Note — If the appliance has both an uploaded certificate and key pair and a generated certificate and key pair, it only uses the certificate and key pair currently selected in the Root Certificate for Signing section.

6. To generate a certificate and key:

   a. Click the Use Generated Certificate and Key option.

   b. Click **Generate New Certificate and Key**.

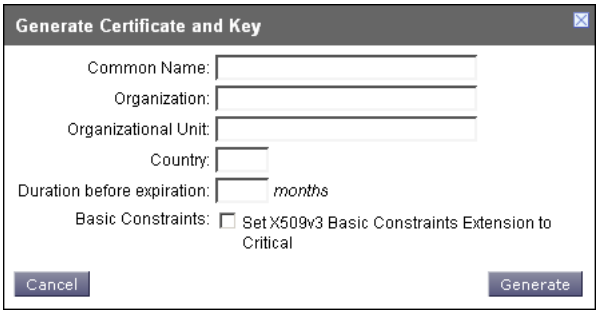

c. In the Generate Certificate and Key dialog box, enter the information to display in the root certificate.

Note — You can enter any ASCII character except the forward slash (/) in the Common Name field.

d. Click **Generate**. The Web Security appliance generates the certificate with the data you entered and generates a key.

The generated certificate information is displayed on the Edit HTTPS Proxy Settings page.

Note — After you generate the certificate and key, you can download the generated certificate to transfer it to the client applications on the network. Do this using the Download Certificate link in the generated key area.

e. Go to step 8 on page 169.

7. To upload a root certificate and key:

a. Click Use Uploaded Certificate and Key.

b. Click **Browse** for the Certificate field to navigate to the certificate file stored on the local machine.

If the file you upload contains multiple certificates or keys, AsyncOS uses the first certificate or key in the file.

Note — The certificate file must be in PEM format. DER format is not supported.

c. Click **Browse** for the Key field to navigate to the private key file. The private key must be unencrypted.

Note — The key length must be 512, 1024, or 2048 bits. Also, the private key file must be in PEM format. DER format is not supported.

d. Click **Upload Files** to transfer the certificate and key files to the Web Security appliance.

The uploaded certificate information is displayed on the Edit HTTPS Proxy Settings page.

Note — After you upload the certificate and key, you can download the certificate to transfer it to the client applications on the network. Do this using the Download Certificate link in the uploaded key area.

8. In the Invalid Certificate Handling section, choose how the appliance handle HTTPS traffic when it encounters invalid server certificates. You can drop, decrypt, or monitor HTTPS traffic for the following types of invalid server certificates:
   - **Expired.** The certificate is either not yet valid, or it is currently past its valid to date.
   - **Mismatched hostname.** The host name in the certificate does not match the host name the client was trying to access. This might happen during a "man in the middle attack," or when a server redirects a request to a different URL. For example, http://mail.google.com gets redirected to http://www.gmail.com.

     Note — AsyncOS can only perform host name match when it is deployed in explicit forward mode. When it is deployed in transparent mode, it does not know the host name of the destination server (it only knows the IP address), so it cannot compare it to the host name in the server certificate.
   - **Unrecognized root authority.** The root certificate authority for the certificate is not in the set of trusted root authorities on the appliance.
   - **All other error types.** Most other error types are due to the appliance not being able to complete the SSL handshake with the HTTPS server. For more information about additional error scenarios for server certificates, see http://www.openssl.org/docs/apps/verify.html.

   Note — When a certificate is both expired and has an unrecognized root authority, the Web Security appliance performs the action specified for an unrecognized root authority.

   For more information about handling invalid server certificates, see "Validating Digital Certificates" on page 159.
9. Submit and commit your changes.

EVALUATING DECRYPTION POLICY GROUP MEMBERSHIP

After AsyncOS assigns an identity to a client request, AsyncOS evaluates the request against the other policy types to determine which policy group it belongs for each type. When HTTPS scanning is enabled, it applies HTTPS requests against the decryption policies. When HTTPS scanning is not enabled, it evaluates HTTP requests against the access policies.

When an HTTPS request gets decrypted, AsyncOS evaluates the decrypted request against the access policies. For more information about how AsyncOS evaluates access policies, see "Evaluating Access Policy Group Membership" on page 122.

AsyncOS applies the configured policy control settings to a client request based on the client request's policy group membership.

To determine the policy group that a client request matches, AsyncOS follows a very specific process for matching the group membership criteria. During this process, it considers the following factors for group membership:

- **Identity.** Each client request either matches an identity or fails authentication and gets terminated. For more information about evaluating identity group membership, see "Evaluating Identity Group Membership" on page 105.
- **Authorized users.** If the assigned identity requires authentication, the user must be in the list of authorized users in the decryption policy group to match the policy group.
- **Advanced options.** You can configure several advanced options for decryption policy group membership. Some of the options (subnet, proxy port, and URL category) can also be defined within the identity. When you configure an advanced option at the decryption policy group level that is also configured within the identity, you narrow down the field of advanced options that the transaction must match.

The information in this section gives an overview of how the appliance matches client requests to decryption policy groups. For more details about exactly how the appliance matches client requests, see "Matching Client Requests to Decryption Policy Groups" on page 170.

AsyncOS sequentially reads through each policy group in the policies table. It compares the client request status to the membership criteria of the first policy group. If they match, AsyncOS applies the policy settings of that policy group.

If they do not match, AsyncOS compares the client request to the next policy group. It continues this process until it matches the client request to a user defined policy group, or if it does not match a user defined policy group, it matches the global policy group. When AsyncOS matches the client request to a policy group or the global policy group, it applies the policy settings of that policy group.

Matching Client Requests to Decryption Policy Groups

Figure 9-7 on page 171 shows how AsyncOS evaluates a client request against the decryption policy groups.

Figure 9-7 Policy Group Flow Diagram for Decryption Policies

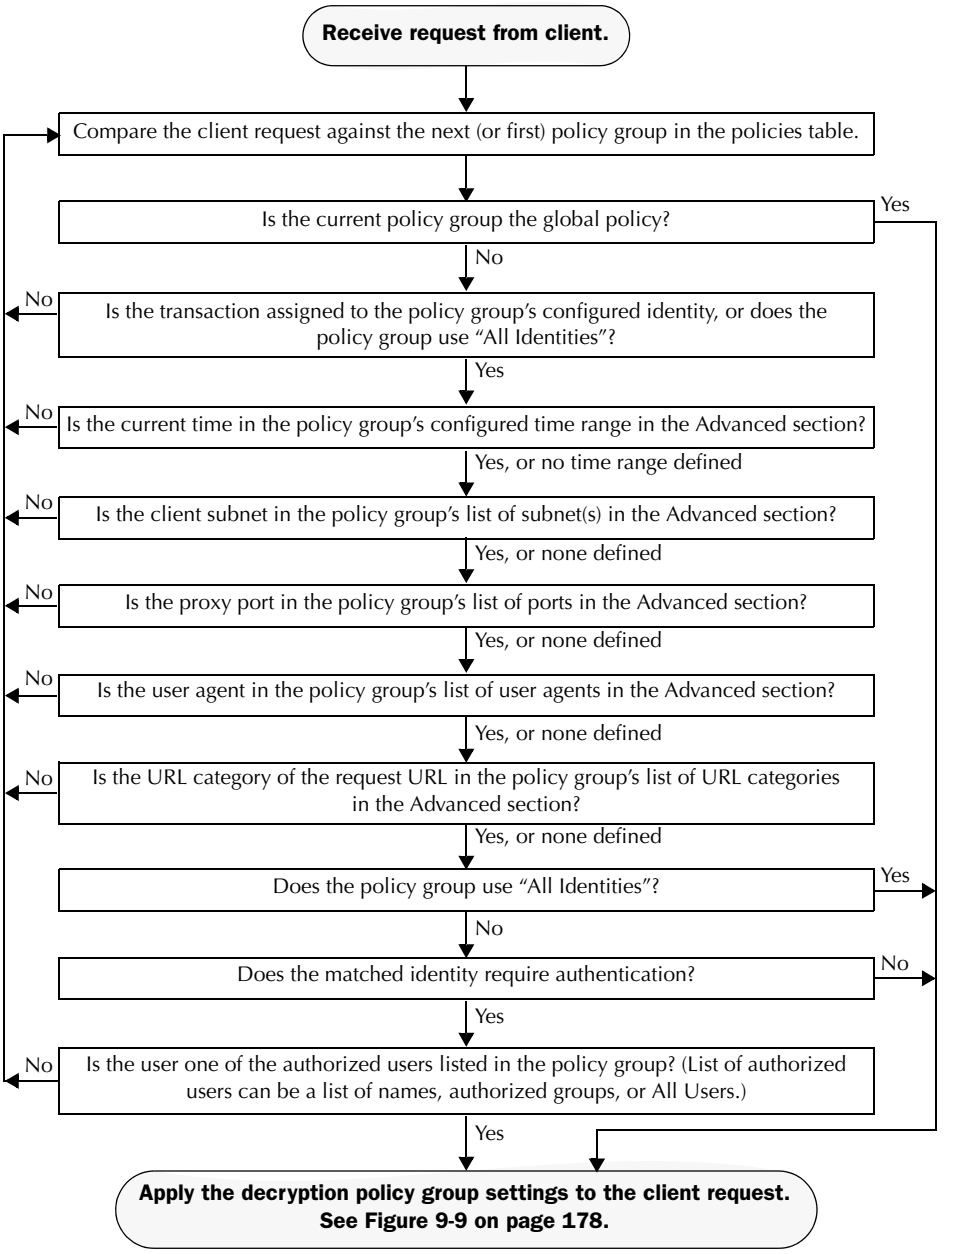

CREATING DECRYPTION POLICIES

You can create decryption policy groups based on combinations of several criteria, such as identity or the URL category of the destination site. You must define at least one criterion for policy group membership. When you define multiple criteria, the client request must meet all criteria to match the policy group.

For more information about how the appliance matches a client request with a policy group, see "Evaluating Decryption Policy Group Membership" on page 170 and "Matching Client Requests to Decryption Policy Groups" on page 170.

You define policy group membership on the Web Security Manager > Decryption Policies page.

To create a decryption policy group:

1. Navigate to the Web Security Manager > Decryption Policies page.
2. Click **Add Policy.**
3. In the Policy Name field, enter a name for the policy group, and in the Description field, optionally add a description.
4. In the Insert Above Policy field, choose where in the policies table to place the policy group.

   When configuring multiple policy groups you must specify a logical order for each group. Carefully order your policy groups to ensure that correct matching occurs.
5. In the Policy Member Definition section, Identity Policy subsection, choose the identity group to apply to this policy group.

   Note — If the identity requires authentication, then authentication information may not be available when a user tries to connect to an HTTPS server. For more information on how HTTPS and authentication work together, see "How Authentication Affects HTTPS Requests" on page 107.
6. If you choose an identity that requires authentication, you can specify which users are authorized for this policy group. These users must authenticate. In the second field in the Identity Policy section, you can choose one of the following options:
   - **All users.** Choose Include all users in this realm.
   - **Specific users.** Choose Specify authorized groups and users, and enter the users in the other fields in the section. The fields that appear depend on the type of authentication used in the identity.
7. Optionally, expand the Advanced section to define additional membership requirements.

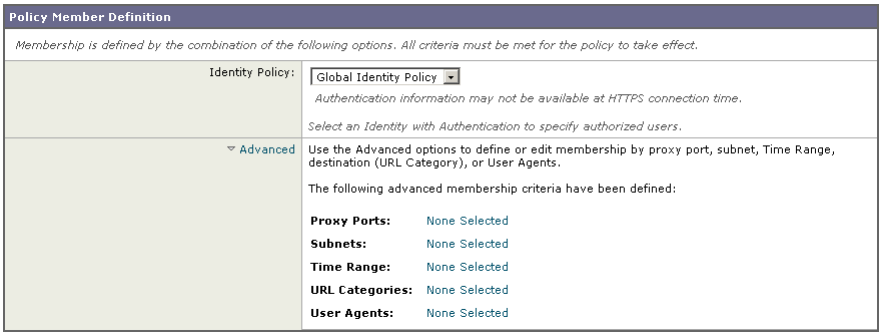

8. To define policy group membership by any of the advanced options, click the link for the advanced option and configure the option on the page that appears.

   Table 9-2 describes the advanced options you can configure for decryption policy groups.

Table 9-2 Decryption Policy Group Advanced Options

Advanced Option	Description
Proxy Ports	Choose whether or not to define policy group membership by the proxy port used to access the Web Proxy. Enter one or more port numbers in the Proxy Ports field. Separate multiple ports with commas. For explicit forward connections, this is the port configured in the browser. For transparent connections, this is the same as the destination port. You might want to define policy group membership on the proxy port if you have one set of clients configured to explicitly forward requests on one port, and another set of clients configured to explicitly forward requests on a different port. IronPort recommends only defining policy group membership by the proxy port when the appliance is deployed in explicit forward mode, or when clients explicitly forward requests to the appliance. When you define policy group membership by the proxy port when clients requests get transparently redirected to the appliance, some requests might be denied. **Note:** The identity associated with this policy group might define identity membership by proxy port. Adding proxy ports in the policy group further narrows down the list of transactions that match this policy group.

Table 9-2 Decryption Policy Group Advanced Options (Continued)

Advanced Option	Description
Subnets	Choose whether or not to define policy group membership by subnet or other addresses. You can choose to use the addresses that may be defined with the associated identity, or you can enter specific addresses here. **Note:** If the associated identity defines identity membership by addresses, you must enter addresses that are a subset of the addresses defined in the identity. Adding addresses in the policy group further narrows down the list of transactions that match this policy group.
Time Range	Choose whether or not to define policy group membership by a defined time range. Choose the time range from the Time Range field and then choose whether this policy group should apply to the times inside or outside the selected time range. For more information on creating time based policies, see "Working with Time Based Policies" on page 93. For more information on creating time ranges, see "Creating Time Ranges" on page 93.
URL Categories	Choose whether or not to define policy group membership by URL categories. Select the user defined or predefined URL categories. **Note:** The identity associated with this policy group might define identity membership by URL categories. Adding URL categories in the policy group further narrows down the list of transactions that match this policy group.
User Agents	Choose whether or not to define policy group membership by the user agent used in the client request. You can select some commonly defined browsers, or define your own using regular expressions. Choose whether this policy group should apply to the selected user agents or to any user agent that is not in the list of selected user agents. For more information on creating user agent based policies, see "Working with User Agent Based Policies" on page 95.

9. Submit your changes.
10. Configure decryption policy group access control settings to define how AsyncOS handles transactions.

    The new policy group automatically inherits global policy group settings until you configure options for each access control component. For more information, see "Controlling HTTPS Traffic" on page 176.

11. Submit and commit your changes.

CONTROLLING HTTPS TRAFFIC

After the Web Security appliance assigns an HTTPS connection request to a decryption policy group, it assigns the access control settings of the policy group to the connection request. The access control settings of the decryption policy group determine whether the appliance allows, drops, or passes through the connection. For more information about the actions the appliance can take on an HTTPS request, see "Decryption Policy Groups" on page 151.

Configure access control settings for decryption policy groups on the Web Security Manager > Decryption Policies page.

Figure 9-8 shows where you can configure access control settings for the decryption policy groups.

Figure 9-8 HTTPS Policies Table

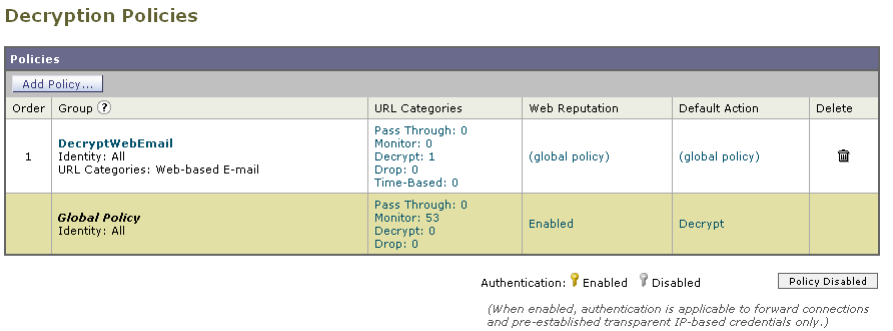

You can configure the following settings to determine what action to take on the HTTPS connection:

- **URL categories.** You can configure the action to take on HTTPS requests for each predefined and custom URL category. Click the link under the URL Categories column for the policy group you want to configure. For more information about working with IronPort URL filters, see "URL Filters" on page 207. For more information about configuring URL categories, see "Configuring URL Filters for Decryption Policy Groups" on page 213.

  Note — If you want to *block* (with end-user notification) a particular URL category for HTTPS requests instead of drop (with no end-user notification), choose to decrypt that URL category in the decryption policy group and then choose to block the same URL category in the access policy group.

- **Web reputation.** You can configure the action to take on HTTPS requests based on the web reputation score of the requested server. Click the link under the Web Reputation column for the policy group you want to configure. For more information about working with web reputation scores, see "Web Reputation in Decryption Policies" on page 232.

- **Default action.** You can configure the action the appliance should take when none of the other settings apply. Click the link under the Default Action column for the policy group you want to configure.

  Note — The configured default action only affects the transaction when no decision is made based on URL category or Web Reputation score. If Web Reputation filtering is disabled, the default action applies to all transactions that match a Monitor action in a URL category. If Web Reputation filtering is enabled, the default action is used only if the Monitor action is selected for sites with no score.

After a decryption policy group is assigned to an HTTPS request, the access control settings for the policy group are evaluated to determine whether to drop, pass through, or decrypt the HTTPS connection request. For more information about assigning a decryption policy group to an HTTPS request, see “Policy Group Membership” on page 90.

Figure 9-9 on page 178 shows how the appliance determines which action to take on an HTTPS request after it has assigned a particular decryption policy to the request.

Figure 9-9 Applying Decryption Policy Actions

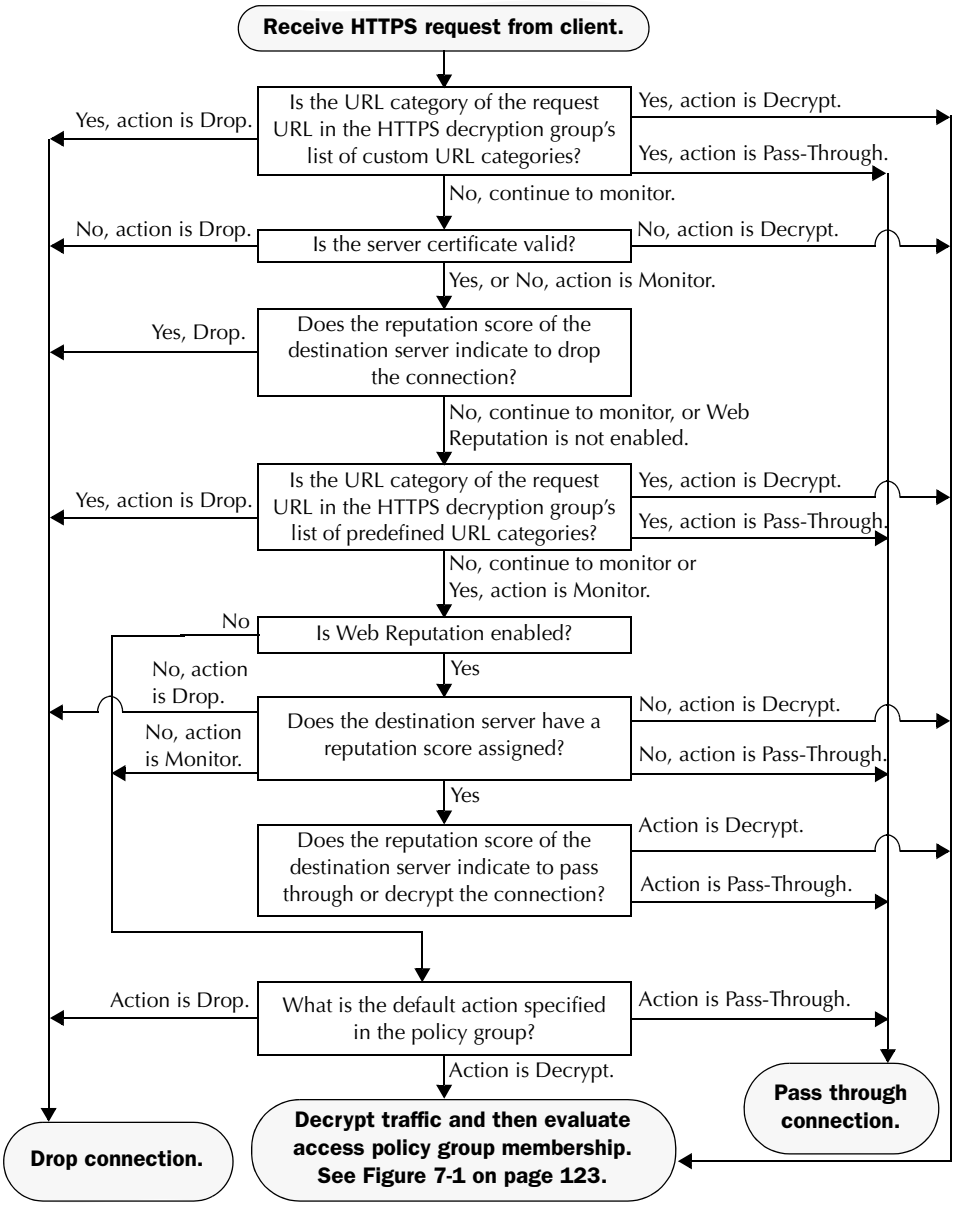

Figure 9-9 shows two different decision points that involve the web reputation score of the destination server. The web reputation score of the server is evaluated only once, but the result is applied at two different points in the decision flow.

For example, note that a web reputation score drop action overrides any action defined for predefined URL categories.

Note — The configured default action only affects the action on the HTTPS request when web reputation filtering is not enabled, or when it is enabled and the server has no score assigned and the action for servers with no scores is to Monitor.

IMPORTING A TRUSTED ROOT CERTIFICATE

When the Web Proxy receives a connection request for an HTTPS server, it validates the trustworthiness of the destination server by verifying the root certificate authority that signed the server certificate. If the Web Proxy does not recognize the root certificate that signed the server certificate, then it does not trust the server certificate. This happens when the HTTPS server uses a certificate authority that is not listed in the set of trusted certificate authorities that ship with the Web Security appliance. This might happen if your organization uses an internal certificate authority to sign certificates for servers on the internal network.

To prevent the Web Proxy from potentially blocking access to servers with unrecognized root certificate authorities, you can upload to the appliance root certificates that your organization trusts. For example, you might want to upload a root certificate used by the servers on your network.

You can upload multiple root certificate files to the appliance, and each file you upload can contain multiple root certificates. However, each certificate you upload must be a root certificate.

To import a trusted root certificate:

1. Navigate to the Security Services > HTTPS Proxy page.

2. In the Custom Root Authority Certificates section, click **Import**.

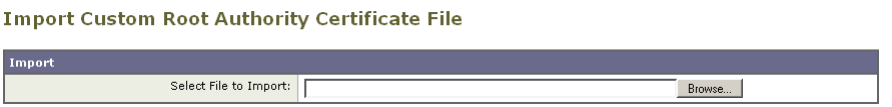

3. In the Import Custom Root Authority Certificate File, click **Browse**.
4. Navigate to the location where the custom root authority certificate file is located and click **Open**.
5. Click **Submit**.

   The uploaded root certificate is displayed in the “Custom Root Authority Certificates” section.
6. Optionally, repeat steps 2 through 5 to upload additional trusted root certificates.
7. Commit your changes.

CHAPTER 10

Notifying End Users

This chapter contains the following information:

NOTIFYING END USERS OF ORGANIZATION POLICIES

The Web Security appliance helps your organization implement and enforce policies for accessing the web. When a policy blocks a user from a website, you can configure the appliance to notify the user why it blocked the URL request. Web users see a webpage that explains that they were blocked from accessing a website and why they were blocked. These pages are called end-user notification pages. AsyncOS can display different end-user notification pages depending on the reason it blocked the URL request. You can use the provided end-user notification pages or define your own.

Configure end-user notification pages on the Security Services > End-User Notification page. Figure 10-1 shows where you configure end-user notification settings.

Figure 10-1 Security Services > End-User Notification Page

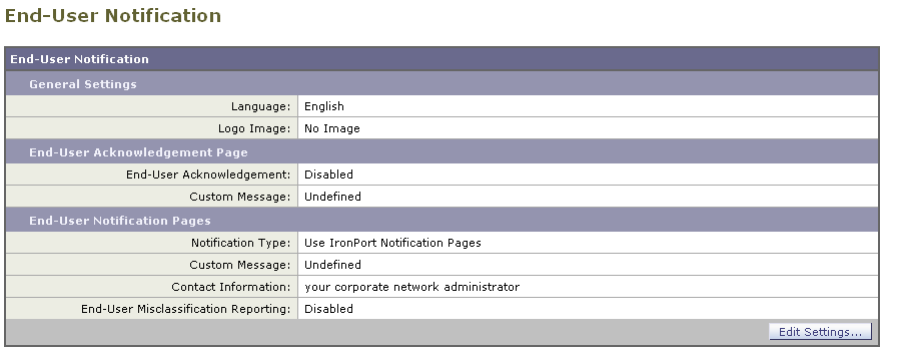

You can configure AsyncOS to display the following types of notification pages:

- **IronPort notification pages.** AsyncOS displays different, predefined notification pages depending on the reason for blocking the URL request. You can customize these pages. For example, you can use your own logo or add custom text. For more information about IronPort notification pages, see “Working With IronPort End-User Notification Pages” on page 185.
- **User defined notification pages.** You can configure AsyncOS to redirect all end-user notification pages to a specific URL. AsyncOS includes parameters in the redirected URL that explain the reasons for the block so the server in the redirected URL can customize the page it displays. For more information about user defined notification pages, see “Working with User Defined End-User Notification Pages” on page 192.
- **End-user acknowledgement page.** You can configure the Web Security appliance to inform users that it is filtering and monitoring their web activity. The appliance does this by displaying an end-user acknowledgement page when a user first accesses a browser after a certain period of time. When the end-user acknowledgement page appears, users must click a link to access the original site requested or any other website. Language and logo settings apply to the end-user acknowledgement page as well as the notification

pages. For more information about configuring the end-user acknowledgement page, see “End-User Acknowledgement Page” on page 194.

CONFIGURING GENERAL SETTINGS FOR NOTIFICATION PAGES

You can configure two general settings for notification pages: logo and language. These general settings apply to all notification pages (acknowledgement, IronPort end-user, and user defined end-user).

To configure the general settings for notification pages:

1. Navigate to the Security Services > End-User Notification page.
2. Click **Edit Settings**.

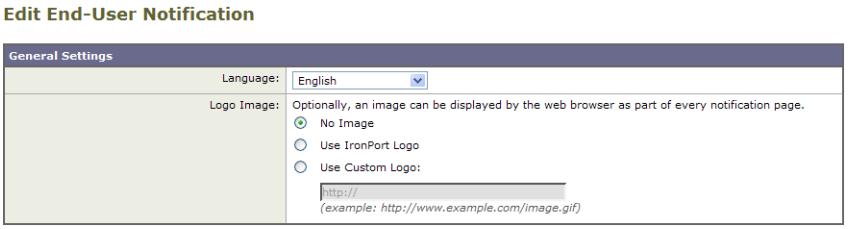

3. Select the language AsyncOS should use when displaying notification pages. You can choose any of the following languages:
   - English
   - French
   - German
   - Italian
   - Spanish
   - Japanese
   - Korean
   - Portuguese
   - Russian
   - Thai
   - Traditional Chinese
   - Simplified Chinese
4. Choose whether or not to use a logo on each notification page. You can specify the IronPort logo or any graphic file referenced at the URL you enter in the Use Custom Logo field.

   Note — See "Custom Text and Logos: Authentication, and End-User Acknowledgement Pages" on page 197 for more information about working with custom logos.

5. Submit and commit your changes.

WORKING WITH IRONPORT END-USER NOTIFICATION PAGES

When you choose end-user notification pages defined by IronPort, AsyncOS displays a different page depending on the reason why it blocked the original page. However, you can still customize each page to make them specific to your organization.

You can customize the following features:

- Custom message
- Contact information
- Allow end-users to report misclassified pages to IronPort

You can also edit the IronPort notification pages stored on the Web Security appliance. For more information about how to do this, see "Editing IronPort Notification Pages" on page 187.

Configuring IronPort Notification Pages

To configure IronPort notification pages:

1. Navigate to the Security Services > End-User Notification page, and click **Edit Settings**.

   The Edit End-User Notification page appears.

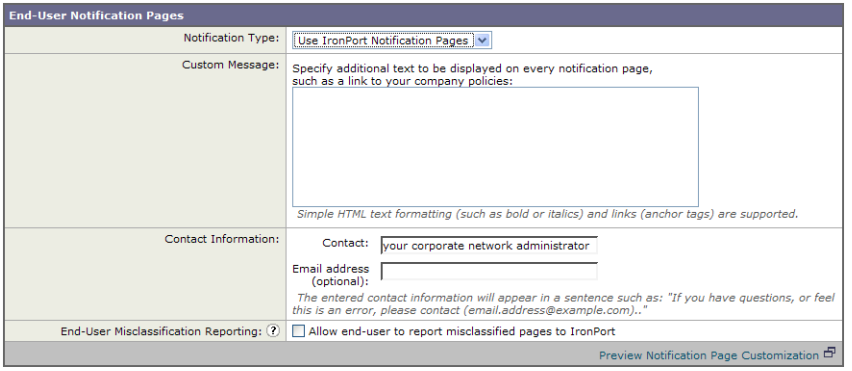

2. From the Notification Type field, choose Use IronPort Notification Pages.
3. Configure the IronPort notification page settings.

Table 10-1 describes the settings you can configure for IronPort notification pages.

Table 10-1 IronPort Notification Page Settings

Setting	Description
Custom Message	Choose whether or not to include additional text you specify on each notification page. When you enter a custom message, AsyncOS places the message before the last sentence on the notification page which includes the contact information. You can include some HTML tags in lower case to format the text. For a list of supported HTML tags, see "Supported HTML Tags in Notification Pages" on page 197. See "Custom Text and Logos: Authentication, and End-User Acknowledgement Pages" on page 197 for more information about working with custom messages.
Contact Information	Choose whether or not to customize the contact information listed on each notification page. AsyncOS displays the contact information sentence as the last sentence on a page, before providing notification codes that users can provide to the network administrator.
End-User Misclassification Reporting	Choose whether or not users can report misclassified URLs to IronPort Systems. When you enable this option, an additional button appears on the IronPort notification pages for sites blocked due to suspected malware or URL filters. This button allows the user to report when they believe the page has been misclassified. It does not appear for pages blocked due to other policy settings. When a user presses this button, data about the blocked request gets sent to the Web Security appliance. AsyncOS logs the information in the Misclassification Log, summarizes the data, and forwards it to IronPort. This feature helps improve efficiency for administrators, and the IronPort Customer Support process. Additionally, misclassification reports improve the efficacy of IronPort URL filtering.

4. Click the "Preview Notification Page Customization" link to view the current end-user notification page in a separate browser window.
5. Submit and commit your changes.

Editing IronPort Notification Pages

Each IronPort Notification page is stored on the Web Security appliance as an HTML file. You can edit the content of these HTML pages to include additional text or to edit the overall look and feel of each page.

To edit the IronPort Notification pages:

1. Use an FTP client to connect to the Web Security appliance.
2. Navigate to the configuration\eun directory.

   In this directory are subdirectories for each supported language for end-user notification pages.
3. Download the language directory files for the IronPort notification pages you want to edit.
4. On your local machine, use a text or HTML editor to edit each HTML file for the IronPort notification pages.

   For a list of rules and guidelines, see "Rules and Guidelines for Editing IronPort Notification Pages" on page 187.
5. Use the FTP client to upload the customized HTML files to the same directory from which you downloaded them in step 3.
6. Open an SSH client and connect to the Web Security appliance.
7. Run the `advancedproxyconfig > miscellaneous` CLI command.
8. Hit Enter until you are prompted with the following question:

   ```
   Enable custom EUN pages?
   ```
9. If the custom end-user notification pages option is currently disabled, type **`1`** to enable it.

   Note — If the custom end-user notification pages option is currently enabled when you update the HTML files, you must first disable it, commit your changes, and then enable it. If you do not do this, the new files do not take effect until the Web Proxy restarts.
10. Commit your change, and close the SSH client.

Rules and Guidelines for Editing IronPort Notification Pages

Use the following rules and guidelines when editing IronPort notification pages:

- Each customized IronPort notification page file must be a valid HTML file. For a list of HTML tags you can include, see "Supported HTML Tags in Notification Pages" on page 197.
- The customized IronPort notification page file names must exactly match the file names shipped with the Web Security appliance.
- Do not include any links to URLs in the HTML files. Any link included in the notification pages are subject to the access control rules defined in the access policies and users might end up in a recursive loop.

- If the configuration\eun directory does not contain a particular file with the required name, then the appliance displays the standard IronPort notification page.
- For new IronPort notification pages to go into effect, you must first upload the customized files to the appliance and then enable the customized files using the `advancedproxyconfig > miscellaneous` CLI command.
- You can use variables in the HTML files to display specific information to the user. Table 10-2 describes the variables you can include in customized end-user notification pages.

Table 10-2 Variables for Customized End-User Notification Pages

Variable	Description
%a	Authentication realm for FTP
%A	ARP address
%b	User-agent name
%B	Blocking reason, such as BLOCK-SRC or BLOCK-TYPE
%c	Error page contact person
%C	Entire Set-Cookie: header line, or empty string
%d	Client IP address
%D	User name
%e	Error page email address
%E	The error page logo URL
%f	User feedback section
%F	The URL for user feedback
%g	The web category name, if available
%G	Maximum file size allowed in MB
%h	The host name of the proxy
%H	The server name of the URL
%i	Transaction ID as a hexadecimal number
%I	Management IP Address
%k	Redirection link for the end-user acknowledgement page

Table 10-2 Variables for Customized End-User Notification Pages (Continued)

Variable	Description
%K	Response file type
%l	WWW-Authenticate: header line
%L	Proxy-Authenticate: header line
%M	The Method of the request, such as “GET” or “POST”
%n	Malware category name, if available
%N	Malware threat name, if available
%p	Proxy connection string
%P	Protocol
%r	Redirect URL
%S	The signature of the proxy
%t	Timestamp in Unix seconds plus milliseconds
%T	The date
%u	The URI part of the URL (the URL excluding the server name)
%U	The full URL of the request
%W	Management WebUI port
%X	Extended blocking code. This is a 16-byte base64 value that encodes the most of the web reputation and anti-malware information logged in the access log, such as the ACL decision tag and WBRS score.
%Y	Admin custom text string, if set, else empty
%%	Prints the percent symbol (%) in the notification page
%?b	This conditional variable evaluates to TRUE if: BLOCK_ADMIN-USER_AGENTS was the deciding ACL rule
%#b	Represents the following condition: `endif` Use this with the %?b conditional variable.

Table 10-2 Variables for Customized End-User Notification Pages (Continued)

Variable	Description
%!b	Represents the following condition: `else` Use this with the %?b conditional variable.
%?e	This conditional variable evaluates to TRUE if: Admin email address is set
%#e	Represents the following condition: `endif` Use this with the %?e conditional variable.
%!e	Represents the following condition: `else` Use this with the %?e conditional variable.
%?f	This conditional variable evaluates to TRUE if: allowUserFeedback is true
%#f	Represents the following condition: `endif` Use this with the %?f conditional variable.
%!f	Represents the following condition: `else` Use this with the %?f conditional variable.
%?N	This conditional variable evaluates to TRUE if: The spyware name is known (non-empty)
%#N	Represents the following condition: `endif` Use this with the %?N conditional variable.
%!N	Represents the following condition: `else` Use this with the %?N conditional variable.
%?Y	This conditional variable evaluates to TRUE if: The Admin custom text string is set

Table 10-2 Variables for Customized End-User Notification Pages (Continued)

Variable	Description
%#Y	Represents the following condition: `endif` Use this with the %?Y conditional variable.
%!Y	Represents the following condition: `else` Use this with the %?Y conditional variable.

WORKING WITH USER DEFINED END-USER NOTIFICATION PAGES

When you choose end-user notification pages defined by someone in your organization, by default, AsyncOS redirects all blocked websites to the URL regardless of the reason why it blocked the original page. However, AsyncOS also passes parameters as a query string appended to the redirect URL so you can ensure that the user sees a unique page explaining the reason for the block.

AsyncOS passes the parameters to the web server as standard URL Parameters in the HTTP GET request. It uses the following format:

```
<notification_page_url>?param1=value1&param2=value2
```

Table 10-3 describes the parameters AsyncOS includes in the query string.

Table 10-3 End-User Notification Parameters for Redirected URLs

Parameter Name	Description
Time	Date and time of the transaction.
ID	Transaction ID.
Client_IP	IP address of the client.
User	Username of the client making the request, if available.
Site	Host name of the destination in the HTTP request.
URI	URL path specified in the HTTP request.
Status_Code	HTTP status code for the request.
Decision_Tag	ACL decision tag as defined in the Access Log entry that indicates how the DVS engine handled the transaction. For more information about ACL decision tags, see “ACL Decision Tags” on page 345.
URL_Cat	URL category that the IronPort URL Filters assigned to the transaction request. Note: AsyncOS sends the entire URL category name for both predefined and user defined URL categories.
WBRS	WBRS score that the Web Reputation Filters assigned to the URL in the request.
DVS_Verdict	Malware category that the DVS engine assigns to the transaction. For more information about malware categories, “Malware Scanning Verdict Values” on page 353.
DVS_ThreatName	The name of the malware found by the DVS engine.

Note — AsyncOS always includes all parameters in each redirected URL. If no value exists for a particular parameter, AsyncOS passes a hyphen (-).

When you want the user to view a different page for each reason for a blocked website, construct a CGI script on the web server that can parse the query string in the redirect URL. Then the server can perform a second redirect to an appropriate page.

Consider the following rules and guidelines when entering the custom URL:

- You can use any HTTP or HTTPS URL.
- The URL may specify a specific port number.
- The URL may not have any arguments after the question mark.
- The URL must contain a well-formed host name.

For example, if you have the following URL entered in the Redirect to Custom URL field:

`http://www.example.com/eun.policy.html`

And you have the following access log entry:

```
1182468145.492 1 172.17.0.8 TCP_DENIED/403 3146 GET http://
www.espn.com/index.html HTTP/1.1 - NONE/- - BLOCK_WEBCAT-DefaultGroup
<Spor,-,-,-,-,-,-,-,-,-,-,-,-> -
```

Then AsyncOS creates the following redirected URL:

```
http://www.example.com/eun.policy.html?Time=21/Jun/
2007:23:22:25%20%2B0000&ID=0000000004&Client_IP=172.17.0.8&User=-
&Site=www.espn.com&URI=index.html&Status_Code=403&Decision_Tag=
BLOCK_WEBCAT-DefaultGroup&URL_Cat=Sports&WBRS=-&DVS_Verdict=-
&DVS_ThreatName=-
```

Configuring User Defined End-User Notification Pages

To configure user defined end-user notification pages:

1. Navigate to the Security Services > End-User Notification page, and click **Edit Settings**.

   The Edit End-User Notification page appears.

2. From the Notification Type field, choose Redirect to Custom URL.
3. In the Notification Page URL field, enter the URL to which you want to redirect blocked websites.

   Note — You can choose whether or not to preview the URL you enter by clicking the Preview Custom URL link.

4. Submit and commit your changes.

END-USER ACKNOWLEDGEMENT PAGE

You can configure the Web Security appliance to inform users that it is filtering and monitoring their web activity. The appliance does this by displaying an end-user acknowledgement page when a user first accesses a browser after a certain period of time. When the end-user acknowledgement page appears, users must click a link to access the original site requested or any other website.

You might want to use an end-user acknowledgement page to force users to explicitly agree to the terms and conditions for browsing the World Wide Web from the organization's network. This might be useful when the Web Proxy is in transparent mode because web users will not otherwise know that their web transactions are being filtered and monitored for security purposes.

When you configure the appliance to display an end-user acknowledgement page, it does so for every user. It displays the end-user acknowledgement page when a user tries to access a website for the first time, or after a configured time interval.

Users are tracked by username if authentication has made a username available, and tracked by IP address if no username is available.

When you enable the end-user acknowledgement page, you can configure the following settings:

- **Maximum time interval.** The time interval determines how often the appliance displays the end-user acknowledgement page for each user. Once a user clicks the link on the end-user acknowledgement page, the appliance considers that user to have acknowledged the proxy for the time you enter for the maximum time interval. This setting applies to users tracked by username and users tracked by IP address. You can specify any value from 30 to 2678400 seconds (one month). Default is 1 day (86400 seconds).
- **Maximum IP address idle timeout.** The maximum IP address idle timeout determines how long a user tracked and acknowledged by IP address can be idle before the user is no longer considered acknowledged. You can specify any value from 30 to 2678400 seconds (one month). Default is four hours (14400 seconds).
- **Custom message.** The custom message is text you enter that appears on every end-user acknowledgement page. You can include some simple HTML tags to format the text. For example, you can change the color and size of the text, or make it italicized. See "Custom Text in Notification Pages" on page 197 for more information.

  Note — You can only include a custom message when you configure the end-user acknowledgement page in the web interface, versus the CLI.

When a user is tracked by IP address, the appliance uses the shortest value for maximum time interval and maximum IP address idle timeout to determine when to display the end-user acknowledgement page again.

Configuring the End-User Acknowledgement Page

You can enable and configure the end-user acknowledgement page in the web interface or the command line interface. However, when you configure the end-user acknowledgement page in the web interface, you can include a custom message that appears on each page. You can include some simple HTML tags in the custom message, such as font color and size.

In the CLI, use `advancedproxyconfig` -> `authentication`.

To configure the end-user acknowledgement page in the web interface:

1. Navigate to the Security Services > End-User Notification page.
2. Click **Edit Settings**.

Figure 10-2 Editing End-User Acknowledgment Page Settings

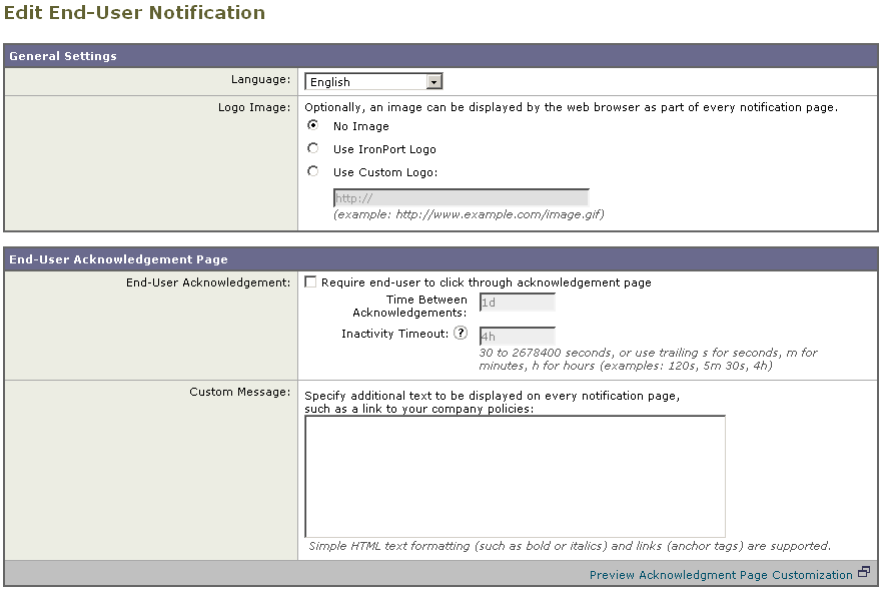

3. In the End-User Acknowledgement Page section, enable the “Require end-user to click through acknowledgement page” field. See “Custom Text and Logos: Authentication, and End-User Acknowledgement Pages” on page 197 for information about how this feature works with custom messages.
4. In the Time Between Acknowledgements field, enter the time interval the appliance uses between displaying the end-user acknowledgement page.

   You can specify any value from 30 to 2678400 seconds (one month). Default is 1 day (86400 seconds). You can enter the value in seconds, minutes, or days. Use ‘s’ for seconds, ‘m’ for minutes, and ‘d’ for days.

5. In the Inactivity Timeout field, enter the maximum IP address idle timeout.

   You can specify any value from 30 to 2678400 seconds (one month). Default is four hours (14400 seconds). You can enter the value in seconds, minutes, or days. Use ‘s’ for seconds, ‘m’ for minutes, and ‘d’ for days.

6. In the Custom Message field, enter any text you want to appear on every end-user acknowledgement page.

   You can include some HTML tags in lower case to format the text. For a list of supported HTML tags, see “Supported HTML Tags in Notification Pages” on page 197.

   For example:

   ```
   Please acknowledge the following statements <i>before</i> accessing
   the Internet.
   ```

7. Click the “Preview Acknowledgment Page Customization” link to view the current end-user acknowledgement page in a separate browser window.

8. Submit and commit your changes.

CUSTOM TEXT IN NOTIFICATION PAGES

The following sections apply to custom text entered for IronPort notification and end-user acknowledgement pages.

Supported HTML Tags in Notification Pages

You can format the text in IronPort notification and end-user acknowledgement pages using some HTML tags. Tags must be in lower case and follow standard HTML syntax (closing tags, etc.).

You can use the following HTML tags.

- <a>
-
-
- <big></big>
-
- <code></code>
-
-
- <small></small>
-

For example, you can make some text italic:

```
Please acknowledge the following statements <i>before</i> accessing the
Internet.
```

With the tag, you can use any CSS style to format text. For example, you can make some text red:

```
<span style="color: red">Warning:</span> You must acknowledge the
following statements <i>before</i> accessing the Internet.
```

Custom Text and Logos: Authentication, and End-User Acknowledgement Pages

All combinations of URL paths and domain names in embedded links within custom text and the custom logo in IronPort notification and end-user acknowledgement pages are exempted from the following:

- User authentication
- End-user acknowledgment
- All scanning, such as malware scanning and URL filters

For example, if the following URLs are embedded in custom text:

`http://www.example.com/index.html`

`http://www.mycompany.com/logo.jpg`

Then all of the following URLs will also be treated as exempt from all scanning:

`http://www.example.com/index.html`

`http://www.mycompany.com/logo.jpg`

`http://www.example.com/logo.jpg`

`http://www.mycompany.com/index.html`

Also, where an embedded URL is of the form: `<protocol>://<domain-name>/ <directory path>/` then all sub-files and sub-directories under that directory path on the host will also be exempted from all scanning.

For example, if the following URL is embedded: `http://www.example.com/gallery2/` URLs such as `http://www.example.com/gallery2/main.php` will also be treated as exempt.

This allows administrators to create a more sophisticated page with embedded content so long as the embedded content is relative to the initial URL. However, you administrators should also take care when deciding which paths to include as links and custom logos.

NOTIFICATION PAGE TYPES

Users accessing the Internet sometimes cannot access the server they want. By default, the Web Security appliance displays a notification page informing users they were blocked and the reason for the block. This section lists and describes all possible notification pages a user might see while accessing the Internet.

Possible reasons that cause notification pages to appear include the following:

- IronPort notification pages are enabled and the user accessed the Internet in a way that violated an access policy.
- IronPort notification pages are configured to allow end-users to report misclassified pages to IronPort and the user reported a misclassified page.
- The end-user acknowledgement page is enabled and the user accessed the Internet for the first time since the timeout period expired.
- HTTPS scanning is enabled and the appliance is configured to drop HTTPS requests to servers with invalid certificates.
- The Web Security appliance could not access the server requested due to an external error, such as DNS failure or an unavailable server.

Most notification pages display a different set of codes that may help administrators or IronPort Customer Support troubleshoot any potential problem. Some codes are for IronPort internal use only.

Table 10-4 describes the different possible codes used in each notification page.

Table 10-4 Codes Used in Notification Pages

Notification Code	Code Description
<version>	The version of the notification message. For IronPort internal use only.
<ACL>	The ACL decision tag that indicates how the DVS engine handled the transaction. For a list of ACL decision tags, see “ACL Decision Tags” on page 345.
<malware_value>	Malware scanning verdict value. For a list of malware scanning verdict values, see “Malware Scanning Verdict Values” on page 353.
<ID>	Transaction ID. For IronPort internal use only.
<time>	The time the error occurred.
<blocking>	Blocking code. For IronPort internal use only.
<HTTP_error>	HTTP error text returned by the web server.

Table 10-4 Codes Used in Notification Pages (Continued)

Notification Code	Code Description
<IP>	Client IP address.
<file_type>	File type of the file the client attempted to download.
<protocol>	The protocol the client requested to use.
<redirected_URL>	The URL to which the client is redirected.
<host_name>	Host name of the web server.

Table 10-5 describes the different notification pages users might encounter.

Table 10-5 Notification Page Types

Notification Title	Notification Text	Notification Codes
Feedback Accepted, Thank You	The misclassification report has been sent. Thank you for your feedback.	N/A
Access Forbidden	Based on your corporate access policies, this web site <*URL*> has been blocked because it has been determined to be a security threat to your computer or the corporate network. Access could also be blocked because this request came from an unrecognized or unauthorized machine.	(<version>, ACCESS_FORBIDDEN, <ACL>, <malware_value>, <ID>, <time>, <blocking>, <HTTP_error>)
Policy: Authentication	Based on your corporate access policies, Internet access has been blocked because the login provided belongs to a user or group that is not allowed Internet access.	(<version>, AUTH, <ACL>)
Bad Request	The system cannot process this request. A non-standard browser may have generated an invalid HTTP request. If you are using a standard browser, please retry the request.	(<version>, BAD_REQUEST)
Policy: Destination	Based on your corporate access policies, access to this web site <*URL*> has been blocked.	(<version>, BLOCK_DEST, <ACL>, <HTTP_error>)
Policy: Source	Based on your corporate access policies, access to this web site <*URL*> has been blocked because this request came from an unauthorized computer.	(<version>, BLOCK_SRC, <ACL>, <IP>, <HTTP_error>)

Table 10-5 Notification Page Types (Continued)

Notification Title	Notification Text	Notification Codes
Security: Browser	Based on your corporate access policies, requests from your computer have been blocked because it has been determined to be a security threat to the corporate network. Your browser may have been compromised by a malware/spyware agent identified as "<*malware name*>". Please contact <*contact name*> <*email address*> and provide the codes shown below. If you are using a non-standard browser and believe it has been misclassified, use the button below to report this misclassification.	(<version>, BROWSER, <ACL>, <browser_type>, <ID>, <time>, <blocking>, <HTTP_error>)
Policy: Browser	Based on your corporate access policies, requests from your browser have been blocked. This browser "<*browser type*>" is not permitted due to potential security risks.	(<version>, BROWSER_CUSTOM, <ACL>, <browser_type>, <ID>, <time>, <HTTP_error>)
Invalid Certificate	A secure session cannot be established because the site <*host name*> provided an invalid certificate.	(<version>, INVALID_CERT, <ACL>, <ID>, <time>, <blocking>, <HTTP_error>)
DNS Failure	The host name resolution (DNS lookup) for this host name <*host name*> has failed. The Internet address may be misspelled or obsolete, the host <*host name*> may be temporarily unavailable, or the DNS server may be unresponsive. Please check the spelling of the Internet address entered. If it is correct, try this request later.	(<version>, DNS_FAIL, <host_name>)
Expectation Failed	The system cannot process the request for this site <*URL*>. A non-standard browser may have generated an invalid HTTP request. If using a standard browser, please retry the request.	(<version>, EXPECTATION_FAILED, <HTTP_error>)
Policy: File Size	Based on your corporate access policies, access to this web site or download <*URL*> has been blocked because the download size exceeds the allowed limit.	(<version>, FILE_SIZE, <ACL>, <HTTP_error>)
Policy: File Type	Based on your corporate access policies, access to this web site or download <*URL*> has been blocked because the file type "<*file type*>" is not allowed.	(<version>, FILE_TYPE, <ACL>, <file_type>, <HTTP_error>)

Table 10-5 Notification Page Types (Continued)

Notification Title	Notification Text	Notification Codes
Filter Failure	The request for page <*URL*> has been denied because an internal server is currently unreachable or overloaded. Please retry the request later.	(<version>, FILTER_FAILURE, <HTTP_error>)
Found	The page <*URL*> is being redirected to <*redirected URL*>.	(<version>, FOUND, <redirected_URL>, <HTTP_error>)
FTP Aborted	The request for the file <*URL*> did not succeed. The FTP server <*host name*> unexpectedly terminated the connection. Please retry the request later.	(<version>, FTP_ABORTED, <HTTP_error>)
FTP Authorization Required	Authentication is required by the FTP server <*host name*>. A valid user ID and password must be entered when prompted. In some cases, the FTP server may limit the number of anonymous connections. If you usually connect to this server as an anonymous user, please try again later.	(<version>, FTP_AUTH_REQUIRED, <host_name>)
FTP Connection Failed	The system cannot communicate with the FTP server <*host name*>. The FTP server may be temporarily or permanently down, or may be unreachable because of network problems. Please check the spelling of the address entered. If it is correct, try this request later.	(<version>, FTP_CONNECTION_FAILED, <host_name>)
FTP Forbidden	Access was denied by the FTP server <*host name*>. Your user ID does not have permission to access this document.	(<version>, FTP_FORBIDDEN, <host_name>)
FTP Not Found	The file <*URL*> could not be found. The address is either incorrect or obsolete.	(<version>, FTP_NOT_FOUND, <HTTP_error>)
FTP Server Error	The system cannot communicate with the FTP server <*host name*>. The FTP server may be temporarily or permanently down, or may not provide this service. Please confirm that this is a valid address. If it is correct, try this request later.	(<version>, FTP_SERVER_ERR, <host_name>)

Table 10-5 Notification Page Types (Continued)

Notification Title	Notification Text	Notification Codes
FTP Service Unavailable	The system cannot communicate with the FTP server <*host name*>. The FTP server may be busy, may be permanently down, or may not provide this service. Please confirm that this is a valid address. If it is correct, try this request later.	(<version>, FTP_SERVICE_UNAVAIL, <host_name>)
Gateway Timeout	The system cannot communicate with the external server <*host name*>. The Internet server may be busy, may be permanently down, or may be unreachable because of network problems. Please check the spelling of the Internet address entered. If it is correct, try this request later.	(<version>, GATEWAY_TIMEOUT, <host_name>)
Internal Error	Internal system error when processing the request for the page <*URL*>. Please retry this request. If this condition persists, please contact <*contact name*> <*email address*> and provide the code shown below.	(<version>, INTERNAL_ERROR, <HTTP_error>)
Security: Malware Risk	Based on your corporate access policies, this web site <*URL*> has been blocked because it has been determined to be a security threat to your computer or the corporate network. This web site has been associated with malware/spyware.	(<version>, MALWARE_GENERAL, <ACL>, <malware_value>, <ID>, <time>, <blocking>, <HTTP_error>)
Security: Malware Detected	Based on your corporate access policies, this web site <*URL*> has been blocked because it has been determined to be a security threat to your computer or the corporate network. Malware <*malware name*> in the category <*malware category*> has been found on this site.	(<version>, MALWARE_SPECIFIC, <ACL>, <malware_value>, <ID>, <time>, <blocking>, <HTTP_error>)
Miss Access Forbidden	This web site <*URL*> has been blocked because it has been determined to be a security threat, based on your corporate access policies.	(<version>, MISS_ACCESS_FORBIDDEN, <HTTP_error>)
No More Forwards	The request for the page <*URL*> failed. The server address <*host name*> may be invalid, or you may need to specify a port number to access this server.	(<version>, NO_MORE_FORWARDS, <HTTP_error>)

Table 10-5 Notification Page Types (Continued)

Notification Title	Notification Text	Notification Codes
Only if Cached But Not in Cache	The page <*URL*> has been blocked based on your corporate policies.	(<version>, ONLY_IF_CACHED_NOT_IN_CACHE, <HTTP_error>)
Policy: General	Based on your corporate access policies, access to this web site <*URL*> has been blocked.	(<version>, POLICY, <ACL>, <ID>, <time>, <blocking>, <HTTP_error>)
Policy: Protocol	Based on your corporate access policies, this request has been blocked because the data transfer protocol "<*protocol type*>" is not allowed.	(<version>, PROTOCOL, <ACL>, <protocol>, <ID>, <time>, <HTTP_error>)
Proxy Authorization Required	Authentication is required to access the Internet using this system. A valid user ID and password must be entered when prompted.	(<version>, PROXY_AUTH_REQUIRED)
Redirect	This request is being redirected. If this page does not automatically redirect, click here to proceed.	N/A
Policy Acknowledgement, Internet Access Policy Acknowledgement	Please acknowledge the following statements before accessing the Internet. Your web transactions will be automatically monitored and processed to detect dangerous content and to enforce corporate policies. By clicking the link below, you acknowledge this monitoring and accept that data about the sites you visit may be recorded. You will be periodically asked to acknowledge the presence of the monitoring system. You are responsible for following corporate polices on Internet access. Click here to accept this statement and access the Internet.	N/A
Proxy Not Licensed	Internet access is not available without proper licensing of the security device. Please contact <*contact name*> <*email address*> and provide the code shown below. Note: To access the management interface of the security device, enter the configured IP address with port.	(<version>, PROXY_UNLICENSED)

Table 10-5 Notification Page Types (Continued)

Notification Title	Notification Text	Notification Codes
Range Not Satisfiable	The system cannot process this request. A non-standard browser may have generated an invalid HTTP request. If you are using a standard browser, please retry the request.	(<version>, RANGE_NOT_SATISFIABLE)
Redirect Permanent	The page <*URL*> is being redirected to <*redirected URL*>.	(<version>, REDIRECT_PERMANENT, <redirected_URL>, <HTTP_error>)
Redirect, Repeat Request	Please repeat your request.	(<version>, REDIRECT_REPEAT_REQUEST)
Server Name Expansion	The server name <*host name*> appears to be an abbreviation, and is being redirected to <*redirected URL*>.	(<version>, SERVER_NAME_EXPANSION, <redirected_URL>, <host_name>)
SOCKS Failure	The server name <*host name*> could not be processed while retrieving the page <*URL*>. This could be due to a problem communicating with the external server.	(<version>, SOCKS_FAIL <HTTP_error>)
URI Too Long	The requested URL was too long and could not be processed. This may represent an attack on your network. Please contact <*contact name*> <*email address*> and provide the code shown below.	(<version>, URI_TOO_LONG)
Policy: URL Filtering	Based on your corporate access policies, access to this web site <*URL*> has been blocked because the web category "<*category type*>" is not allowed.	(<version>, WEBCAT, <ACL>, <ID>, <time>, <blocking>, <HTTP_error>)
WWW Authorization Required	Authentication is required to access the requested web site <*host name*>. A valid user ID and password must be entered when prompted.	(<version>, WWW_AUTH_REQUIRED, <host_name>)

CHAPTER 11

URL Filters

This chapter contains the following information:

URL FILTERS OVERVIEW

IronPort URL Filters allow administrators to control user access based on the web server category of a particular HTTP or HTTPS request. For example, you can block all HTTP requests for gambling web sites, or you can decrypt all HTTPS requests for web-based email websites.

Using policy groups, you can create secure policies that control access to web sites containing objectionable or questionable content. The sites that are actually blocked, dropped, allowed, or decrypted depend on the categories you select when setting up category blocking for each policy group.

You can use URL filters to perform the following tasks:

- **Define policy group membership.** You can define policy group membership by the URL category of the request URL.
- **Control access to HTTP and HTTPS requests.** You can choose to allow or block HTTP requests by URL category using access policy groups, and you can choose to pass through, drop, or decrypt HTTPS requests by URL category using decryption policy groups. For more information, "Configuring IronPort URL Filters" on page 211.

The Web Security appliance ships with over 50 predefined URL categories, such as Arts, Education, Hacking, Web-based Email, and more. However, you can also create user defined custom URL categories that specify specific host names and IP addresses. For more information about working with custom URL categories, see "Custom URL Categories" on page 216.

Matching URLs to URL Categories

When the IronPort URL Filters match a URL category to the URL in a client request, it first evaluates the URL against the custom URL categories. If the URL in the request does not match a custom category, the IronPort URL Filters compares it to the predefined URL categories. If the URL does not match any custom or predefined URL categories, the request is uncategorized. For more information about uncategorized URLs, see "Uncategorized URLs" on page 209.

The IronPort URL Filters Database

The Web Security appliance collects information and maintains its own filtering categories database. The filtering categories database periodically receives updates from the IronPort update server (`https://update-manifests.ironport.com`). Server updates are automated, and the update interval is set by the server as opposed to the appliance. Updates to the database occur with a regular degree of frequency, and require no administrator intervention.

For information about update intervals and the IronPort update server, see "Component Updates" on page 391.

Uncategorized URLs

An uncategorized URL applies to any request that does not match the context of any pre-defined or custom URL category. All transactions resulting in unmatched categories are reported on the Monitor > URL Categories page as "Uncategorized URLs." A large number of uncategorized URLs are generated from requests to web sites within the internal network. Because this type of internal transaction can falsely inflate reporting data and misrepresent the efficacy of IronPort URL Filters, IronPort recommends using custom URL categories to group internal URLs and allow all requests to internal web sites. This decreases the number of web transactions reported as "Uncategorized URLs" and instead reports internal transactions as part of "URL Filtering Bypassed" statistics.

For more information, see "Understanding Unfiltered and Uncategorized Data" on page 222.

For more information about creating custom URL categories, see "Custom URL Categories" on page 216.

ENABLING IRONPORT URL FILTERS

To apply predefined category settings to policy groups and configure custom settings to manage web transactions, you must first enable IronPort URL Filters. By default, IronPort URL Filters are enabled in the System Setup Wizard.

When you enable this feature, you can choose the default action AsyncOS should use when the IronPort URL Filters service is unavailable, either monitor or block.

Note — When the IronPort URL Filters service is unavailable and this setting is set to block, HTTP transactions are blocked, but HTTPS transaction are allowed to proceed to the next decision. For more information, see Figure 9-9 on page 178.

To enable IronPort URL Filters and modify the default action:

1. Navigate to the Security Services > IronPort URL Filters page.
2. Click **Edit Global Settings**.

   The Edit IronPort URL Filters Settings page appears.

3. Verify the Enable IronPort URL Filters property is enabled.
4. Choose the default action AsyncOS should use when the IronPort URL Filters service is unavailable, either Monitor or Block. Default is Monitor.

   Note — The default action only applies for access policies, not decryption policies.

5. Submit and commit your changes.

CONFIGURING IRONPORT URL FILTERS

IronPort URL Filters allow you to filter HTTP requests in access policies, and HTTPS requests in decryption policies. To configure URL filtering in a policy group, click the link in the policies table under the URL Categories column for the policy group you want to edit. For more information about the policies table, see "Using the Policies Tables" on page 87.

When you configure URL categories for policy groups, the following sections are displayed where you can configure URL filtering:

- **Custom URL Category Filtering.** You can configure filtering for a custom category when a custom category is configured on the Web Security Manager > Custom URL Categories page and applies to that policy group type. For more information about custom URL categories, see "Custom URL Categories" on page 216.
- **Predefined URL Category Filtering.** The Web Security appliance ships with a set of predefined URL categories.

The URL filtering actions you can configure depends on the type of policy group.

- **Access policies.** See "Configuring URL Filters for Access Policy Groups" on page 211.
- **Decryption policies.** See "Configuring URL Filters for Decryption Policy Groups" on page 213.

Configuring URL Filters for Access Policy Groups

You can configure URL filtering for user defined access policy groups and the Global Policy Group.

To configure URL filtering in an access policy group:

1. Navigate to the Web Security Manager > Access Policies page.
2. Click the link in the policies table under the URL Categories column for the policy group you want to edit.

   The Access Policies: URL Categories: *policyname* page appears.

Figure 11-1 Configuring Access Policy URL Categories

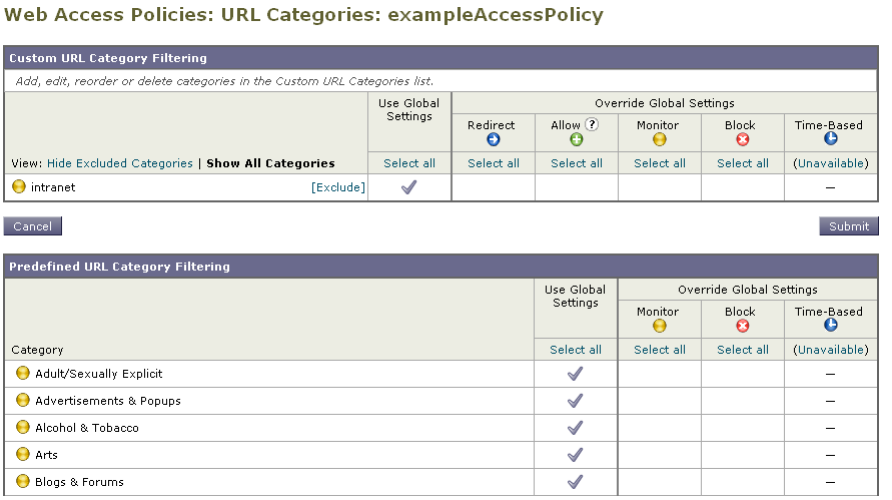

3. In the Custom URL Category Filtering section, choose an action for each custom URL category. Table 11-1 describes each action.

Table 11-1 URL Category Filtering for Access Policies

Action	Description
Exclude (Include)	Choose whether or not the IronPort URL Filters should compare the client request against the custom URL category. Excluding a custom URL category ignores the category completely, acting as if it does not exist in the first place. By default, IronPort URL Filters compares the URL in a client request to custom URL categories before predefined URL categories. You can override this behavior by clicking the Exclude link for a particular custom URL category. This enables you to configure the access policy to ignore the custom URL category without deleting it. Excluding a custom category is different than "Use Global Settings" which uses the settings for the same URL category as defined in the global policy. The Exclude link ignores the custom category settings and instead uses the predefined URL category settings (or another custom URL category, if applicable) for that URL. Once you click the Exclude link, it changes to an Include link to allow you to include the category in the policy again.
Use Global Setting	Uses the action for this category in the Global Policy Group. This is the default action for user defined policy groups. Applies to user defined policy groups only.

Table 11-1 URL Category Filtering for Access Policies (Continued)

Action	Description
Redirect	Redirects traffic originally destined for a URL in this category to a location you specify. When you choose this action, the Redirect To field appears. Enter a URL to which to redirect all traffic. For more information about redirecting traffic, see "Redirecting Traffic" on page 219.
Allow	Always allows client requests for web sites in this category. Allowed requests bypass all further filtering and malware scanning. Only use this setting for trusted web sites. You might want to use this setting for internal sites.
Monitor	The Web Proxy neither allows nor blocks the request. Instead, it continues to evaluate the client request against other policy group access control settings, such as web reputation filtering.
Block	The Web Proxy denies transactions that match this setting.
Time-Based	The Web Proxy blocks or monitors the request during the time ranges you specify. For more information about creating time based URL filtering actions, see "Creating Time Based URL Filters" on page 221.

4. In the Predefined URL Category Filtering section, choose one of the following actions for each category:
   - Exclude link
   - Use Global Settings
   - Monitor
   - Block
   - Time-Based

   See Table 11-1 for details on these actions.
5. In the Uncategorized URLs section, choose the action to take for client requests to web sites that do not fall into a predefined or custom URL category.
6. Submit and commit your changes.

Configuring URL Filters for Decryption Policy Groups

You can configure URL filtering for user defined decryption policy groups and the global decryption policy group.

To configure URL filtering in a decryption policy group:

1. Navigate to the Web Security Manager > Decryption Policies page.
2. Click the link in the policies table under the URL Categories column for the policy group you want to edit.

   The Decryption Policies: URL Categories: *policyname* page appears.

   Figure 11-2 Configuring Decryption Policy URL Categories

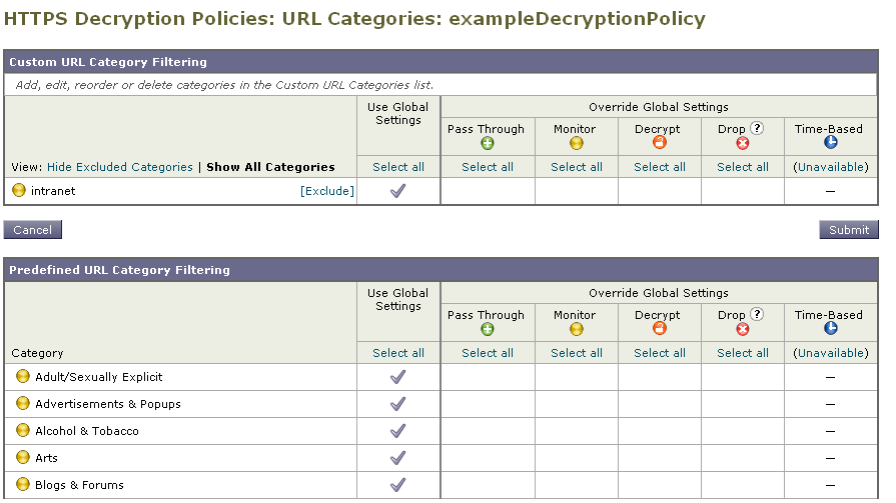

3. Choose an action for each custom and predefined URL category. Table 11-2 describes each action.

   Table 11-2 URL Category Filtering for Decryption Policies

Action	Description
Exclude (Include)	Choose whether or not the IronPort URL Filters should compare the client request against the custom URL category. By default, IronPort URL Filters compares the URL in a client request to custom URL categories before predefined URL categories. You can override this behavior by clicking the Exclude link for a particular custom URL category. This enables you to configure the access policy to ignore the custom URL category without deleting it. Once you click the Exclude link, it changes to an Include link to allow you to include the category in the policy again. This option only applies to custom URL categories.
Use Global Setting	Uses the action for this category in the global decryption policy group. This is the default action for user defined policy groups. Applies to user defined policy groups only.

Table 11-2 URL Category Filtering for Decryption Policies (Continued)

Action	Description
Pass Through	Passes through the connection between the client and the server without inspecting the traffic content. You might want to pass through connections to trusted secure sites, such as well known banking and financial institutions.
Monitor	The Web Proxy neither allows nor blocks the request. Instead, it continues to evaluate the client request against other policy group access control settings, such as web reputation filtering.
Decrypt	Allows the connection, but inspects the traffic content. The appliance decrypts the traffic and applies access policies to the decrypted traffic as if it were a plaintext HTTP connection. By decrypting the connection and applying access policies, you can scan the traffic for malware. You might want to decrypt connections to third party email providers, such as gmail or hotmail. For more information about how the appliance decrypts HTTPS traffic, see "Decrypting HTTPS Traffic" on page 160.
Drop	Drops the connection and does not pass the connection request to the server. The appliance does not notify the user that it dropped the connection. You might want to drop connections to third party proxies that allow users on the network bypass the organization's acceptable use policies.

Note — If you want to *block* a particular URL category for HTTPS requests, choose to decrypt that URL category in the decryption policy group and then choose to block the same URL category in the access policy group.

4. In the Uncategorized URLs section, choose the action to take for client requests to web sites that do not fall into a predefined or custom URL category. You can choose any action listed in Table 11-2.
5. Submit and commit your changes.

CUSTOM URL CATEGORIES

The Web Security appliance ships with over 50 predefined URL categories, such as Arts, Education, Hacking, Web-based Email, and more. However, you can also create user defined custom URL categories that specify specific host names and IP addresses. You might want to create custom URL categories for internal sites or a group of external sites you know you can trust.

Create, edit, and delete custom URL categories on the Web Security Manager > Custom URL Categories page.

Figure 11-3 Custom URL Categories Page

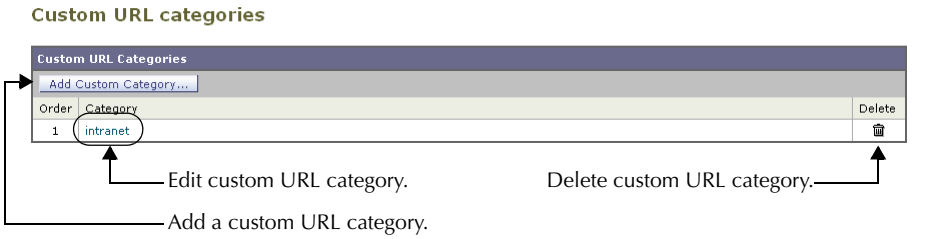

Note — The Web Security appliance uses the first four characters of custom URL category names preceded by "c_" in the access logs. Consider the custom URL category name if you use Sawmill for IronPort to parse the access logs. If the first four characters of the custom URL category include a space, Sawmill for IronPort cannot properly parse the access log entry. Instead, only use supported characters in the first four characters if you will use Sawmill for IronPort to parse the access logs. If you want to include the full name of a custom URL category in the access logs, add the %XF format specifier to the access logs. For more information on how to do this, see "Custom Formatting" on page 356.

To create or edit a custom URL category:

1. Navigate to the Web Security Manager > Custom URL Categories page.
2. To create a custom URL category, click **Add Custom Category**. To edit an existing custom URL category, click the name of the URL category.

Figure 11-4 Creating a Custom URL Category

3. Enter the settings in Table 11-3 for the custom URL category.

Table 11-3 Custom URL Category Settings

Setting	Description
Category Name	Enter a name for the URL category. This name appears when you configure URL filtering for policy groups.
List Order	Choose the order in the list of custom URL categories to place this category. Enter "1" for the topmost URL category. IronPort URL Filters evaluate a client request against the custom URL categories in the order specified.
Sites	Enter one or more addresses that belong in the custom category. You can enter multiple addresses separated by line breaks or commas. You can enter addresses using any of the following formats: • IP address, such as 10.1.1.0 • CIDR address, such as 10.1.1.0/24 • Domain name, such as example.com • Hostname, such as crm.example.com • Partial hostname, such as .example.com **Note:** Entering a partial hostname, such as .example.com, also matches www.example.com.

Table 11-3 Custom URL Category Settings (Continued)

Setting	Description
Advanced: Regular Expressions	You can use regular expressions to specify multiple web servers that match the pattern you enter. **Note:** IronPort URL Filters compares URLs with addresses entered in the Sites field first. If the URL of a transaction matches an entry in the Sites field, it is not compared to any expression entered here. For more information about using regular expressions in the Web Security appliance, see "Regular Expressions" on page 223.

4. Submit and commit your changes.

REDIRECTING TRAFFIC

In addition to using the Web Security appliance to monitor and block traffic to certain websites, you can also use it to redirect users to a different website. You can configure the appliance to redirect traffic originally destined for a URL in a custom URL category to a location you specify. This allows you to redirect traffic at the appliance instead of at the destination server.

You might want to redirect traffic at the appliance if your organization published the links to an internal site, but the location of the site changed since publication, or if you do not have control over the web server.

Configure the appliance to redirect custom URL categories to another location when you configure the URL categories for an access policy group. You can redirect traffic for a custom access policy group or the Global Policy Group.

Note — To redirect traffic, you must define at least one custom URL category. For more information about creating custom URL categories, see "Custom URL Categories" on page 216.

Note — Beware of infinite loops when you configure the appliance to redirect traffic. For example, if you redirect traffic destined for http://A.example.com to http://B.example.com and you also inadvertently redirect traffic destined for http://B.example.com to http://A.example.com, then you create an infinite loop. In this case, the appliance redirects the traffic back and forth between the two URLs indefinitely.

To redirect traffic:

1. Navigate to the Web Security Manager > Access Policies page.
2. Click the link under the URL Categories column for an access policy group or the Global Policy Group.

   The Access Policies: URL Categories: *policyname* page appears.
3. In the Custom URL Category Filtering section, click the Redirect column for the custom category you want to redirect.
4. Enter the URL to which you want to redirect traffic in the Redirect To field for the custom category.

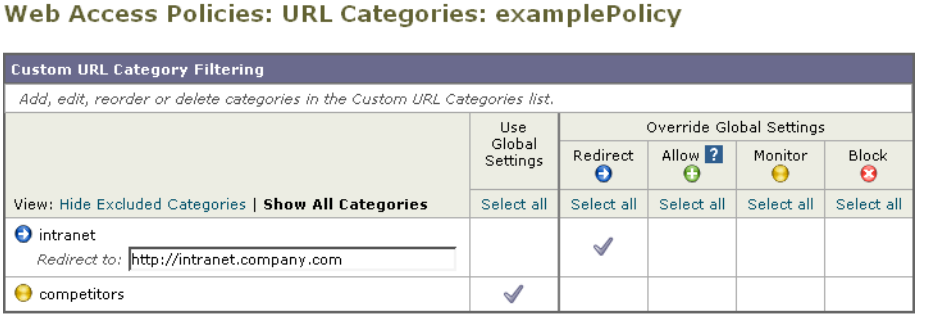

5. Submit and commit your changes.

CREATING TIME BASED URL FILTERS

You can configure how the Web Security appliance to handles requests for URLs in particular categories differently based on time and day. For example, you can block access to social networking sites, such as blogs and forums, during business hours.

To define URL filtering actions by time you must first define at least one time range. For information about time ranges, see "Working with Time Based Policies" on page 93.

To create time based URL filtering actions for an access policy:

1. Navigate to the Web Security Manager > Access Policies page.
2. Click the link in the policies table under the URL Categories column for the policy group you want to edit.

   The Access Policies: URL Categories: *policyname* page appears.
3. Select Time-Based for the custom or predefined URL category you want to configure based on time range.

   Figure 11-5 Defining Time Based URL Filtering Actions

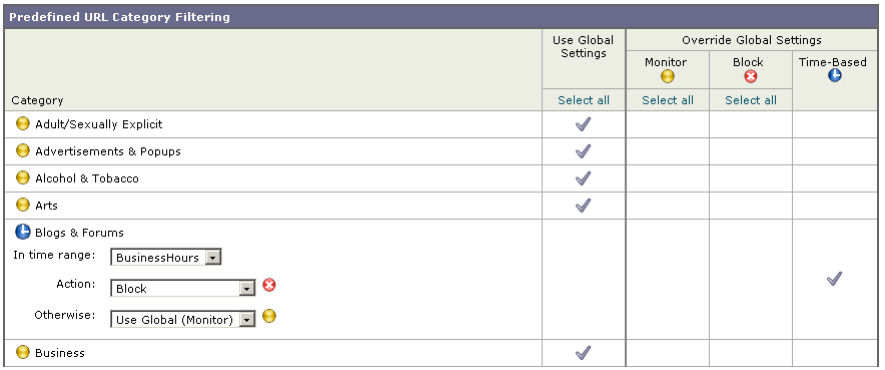

   When you select Time-Based for the URL category, additional fields appear under the category name where you can choose the actions.
4. In the In Time Range field, choose the defined time range to use for the URL category.

   For information about defining time ranges, see "Creating Time Ranges" on page 93.
5. In the Action field, choose the action to enact on transactions in this URL category during the defined time range.
6. In the Otherwise field, choose the action to enact on transactions in this URL category *outside* the defined time range.
7. Submit and commit your changes.

VIEWING URL FILTERING ACTIVITY

The Monitor > URL Categories page provides a collective display of URL statistics that includes information about top URL categories matched and top URL categories blocked. Additionally, this page displays category-specific data for bandwidth savings and web transactions. For detailed information about monitoring and reporting functionality, see "Monitoring" on page 305.

Understanding Unfiltered and Uncategorized Data

When viewing URL statistics on the Monitor > URL Categories page, it is important to understand how to interpret the following data:

- **URL Filtering Bypassed** — This data represents policy, port, and admin user agent blocking that occurs before URL filtering.
- **Uncategorized URL** — This data represents all transactions for which the filtering engine is queried, but no category is matched.

Access Log File

The access log file records the URL category for each transaction in the Web Reputation filtering and anti-malware scanning section of each entry. For more information about the access log, see "Access Log File" on page 343. For a list of each URL category, see "URL Category Abbreviations" on page 350.

REGULAR EXPRESSIONS

Regular expressions are pattern matching descriptions that contain normal printable characters and special characters that are used to match patterns in text strings. For example, a text string such as "welcome" matches "welcome" or "welcomemyfriend." When a match occurs, the function returns true. If no match occurs, the function returns false. Actions are executed only when a pattern-matching expression is true.

The Web Security appliance uses POSIX extended regular expression syntax, fully described by IEEE POSIX 1003.2. However, the appliance does not support using a backward slash to escape a forward slash. If you need to use a forward slash in a regular expression, type the forward slash without a backward slash.

Note — Technically, AsyncOS for Web uses the Flex regular expression analyzer. For more detailed information about how it reads regular expressions, see http://flex.sourceforge.net/manual/Patterns.html.

You can use regular expressions in the following locations:

- **Custom URL categories for access policies.** When you create a custom URL category to use with access policy groups, you can use regular expressions to specify multiple web servers that match the pattern you enter. For more information about creating custom URL categories, see "Custom URL Categories" on page 216.
- **Custom user agents to block.** When you edit the applications to block for an access policy group, you can use regular expressions to enter specific user agents to block, such as Skype or Microsoft Internet Explorer. For more information about using regular expressions to block user agents, see "Policy: Applications" on page 133.

Note — Regular expressions that perform extensive character matching consume resources and can affect system performance. For this reason, regular expressions should be cautiously applied.

Forming Regular Expressions

Regular expressions are rules that typically use the word "matches" in the expression. They can be applied to match specific URL destinations or web servers. For example, the following regular expression matches any pattern containing blocksite.com:

```
\.blocksite\.com
```

Consider the following regular expression example:

```
server[0-9]\.example\.com
```

In this example, server[`0-9`] matches `server0`, `server1`, `server2`, ..., `server9` in the domain `example.com`.

In the following example, the regular expression matches files ending in `.exe`, `.zip`, and .`bin` in the `downloads` directory.

```
/downloads/.*\.(exe|zip|bin)
```

Avoid using regular expressions strings that are redundant because they can cause higher CPU usage on the Web Security appliance. A redundant regular expression is one that starts or ends with ".*".

Note — You must enclose regular expressions that contain blank spaces or non-alphanumeric characters in ASCII quotation marks.

Regular Expression Character Table

Table 11-4 describes characters that are commonly used to form regular expressions:

Table 11-4 Regular Expression Character Descriptions

Character	Description
.	Matches a single character.
*	Matches zero or more occurrences of the preceding regular expression. For example: [0-9]* matches any number of digits ".*" matches any arbitrary string of characters
^	Matches the beginning of a line as the first character of a regular expression.
$	Matches the end of a line as the last character of a regular expression.
+	Matches one or more occurrences of the preceding regular expression.
?	Matches zero or one occurrence of the preceding regular expression.
|	Matches the preceding regular expression or the following regular expression. For example: x|y matches either x or y abc|xyz matches either of the strings abc or xyz
[ ]	Matches the characters or digits that are enclosed within the brackets. For example: [a-z] matches any character between a and z [r-u] matches any of the characters r, s, t, or u [0-3] matches any of the single digits 0, 1, 2, 3
{ }	Specifies the number of times to match the previous pattern. For example: D{1,3} matches one to three occurrences of the letter D

Table 11-4 Regular Expression Character Descriptions (Continued)

Character	Description
( )	Group characters in a regular expression. For example: (abc)* matches abc or abcabcabc
"..."	Literally interprets any characters enclosed within the quotation marks.
\	Escape character.

Note — To match the literal version of any of the special characters, the character must be preceded by a backslash "\". For example, to exactly match a period "." the regular expression must use "\." as in "\.example\.com". However, the appliance does not support using a backward slash to escape a forward slash. If you need to use a forward slash in a regular expression, type the forward slash without a backward slash.

CHAPTER 12

Web Reputation Filters

This chapter contains the following information:

WEB REPUTATION FILTERS OVERVIEW

IronPort Web Reputation Filters is a security feature that analyzes web server behavior and assigns 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 web access and decryption policies.

The Web Reputation Database

The Web Security appliance collects information and maintains a filtering database that contains aggregated traffic statistics, request attributes, and information about how different types of requests are handled. Additionally, the appliance can be configured to send web reputation statistics to a SenderBase server. SenderBase server information is leveraged with data feeds from the IronPort Common Security Database (SenderBase® Network) and the collective information is used to produce a Web Reputation Score.

Note — For more information, see "The SenderBase Network" on page 4.

Maintaining the Database Tables

The web reputation filtering component periodically receives updates to its database tables from the IronPort update server (`https://update-manifests.ironport.com`). Server updates are automated, and the update interval is set by the server as opposed to the appliance. Updates to the database tables occur with a regular degree of frequency, and require no administrator intervention.

For information about update intervals and the IronPort update server, see "Component Updates" on page 391.

WEB REPUTATION SCORES

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

Note — IronPort does not collect personally identifiable information such as usernames, passwords, or client IP addresses.

ENABLING WEB REPUTATION FILTERS

To use web reputation in policy groups, you must first enable Web Reputation Filters. By default, Web Reputation Filters are enabled in the System Setup Wizard. If it is not enabled in the System Setup Wizard, you can edit them in the web interface.

To enable Web Reputation Filters in the web interface:

1. Navigate to the Security Services > Web Reputation Filters page.
2. Click **Enable**.

   The Web Reputation Filters License Agreement appears.
3. Read the terms of the Web Reputation Filters License Agreement, and click **Accept**.
4. Click **Edit Settings**.

   The Edit Web Reputation Filters Settings page appears.
5. Verify the Enable Web Reputation Filters property is enabled.
6. Submit and commit your changes.

HOW WEB REPUTATION FILTERING WORKS

Web Reputation Scores are associated with an action to take on a URL request. The available actions depend on the policy group type that is assigned to the URL request:

- **Access policies.** You can choose to block, scan, or allow.
- **Decryption policies.** You can choose to drop, decrypt, or pass through.

You can configure each policy group to correlate an action to a particular Web Reputation Score.

Web Reputation in Access Policies

Table 12-1 describes the default Web Reputation Scores for access policies.

Table 12-1 Default Web Reputation Scores for Access Policies

Score	Action	Description	Example
-10 to -6.0	Block	Bad site. The request is blocked, and no further malware scanning occurs.	• URL downloads information without user permission. • Sudden spike in URL volume. • URL is a typo of a popular domain.
-5.9 to 5.9	Scan	Undetermined site. Request is passed to the DVS engine for further malware scanning. The DVS engine scans the request and server response content.	• Recently created URL that has a dynamic IP address and contains downloadable content. • Network owner IP address that has a positive Web Reputation Score.
6.0 to 10.0	Allow	Good site. Request is allowed. No malware scanning required.	• URL contains no downloadable content. • Reputable, high-volume domain with long history. • Domain present on several allow lists. • No links to URLs with poor reputations.

For example, by default, URLs in an HTTP request that are assigned a Web Reputation Score of +7 are allowed and require no further scanning. However, a weaker score for an HTTP request, such as +3, is automatically forwarded to the IronPort DVS engine where it is scanned for malware. Any URL in an HTTP request that has a very poor reputation is blocked.

Web Reputation in Decryption Policies

Table 12-2 describes the default Web Reputation Scores for access policies.

Table 12-2 Default Web Reputation Scores for Decryption Policies

Score	Action	Description
-10 to -9.0	Drop	Bad site. The request is dropped with no notice sent to the end user. Use this setting with caution.
-8.9 to 5.9	Decrypt	Undetermined site. Request is allowed, but the connection is decrypted and the decrypted traffic is applied to access policies. For more information about how the appliance decrypts HTTPS traffic, see “Decrypting HTTPS Traffic” on page 160.
6.0 to 10.0	Pass through	Good site. Request is passed through with no inspection or decryption.

CONFIGURING WEB REPUTATION SCORES

When you install and set up the Web Security appliance, it has default settings for Web Reputation Scores. However, you can modify threshold settings for web reputation scoring to fit your organization's needs.

You configure the web reputation filter settings for each policy group.

Configuring Web Reputation for Access Policies

To edit the web reputation filter settings for an access policy group:

1. Navigate to the Web Security Manager > Access Policies page.
2. Click the link under the Web Reputation and Anti-Malware Filtering column for the access policy group you want to edit.
3. Under the Web Reputation and Anti-Malware Settings section, choose "Define Web Reputation and Anti-Malware Custom Settings" from the drop down menu if it is not selected already.

Figure 12-1 Web Reputation Filter Settings for Access Policies

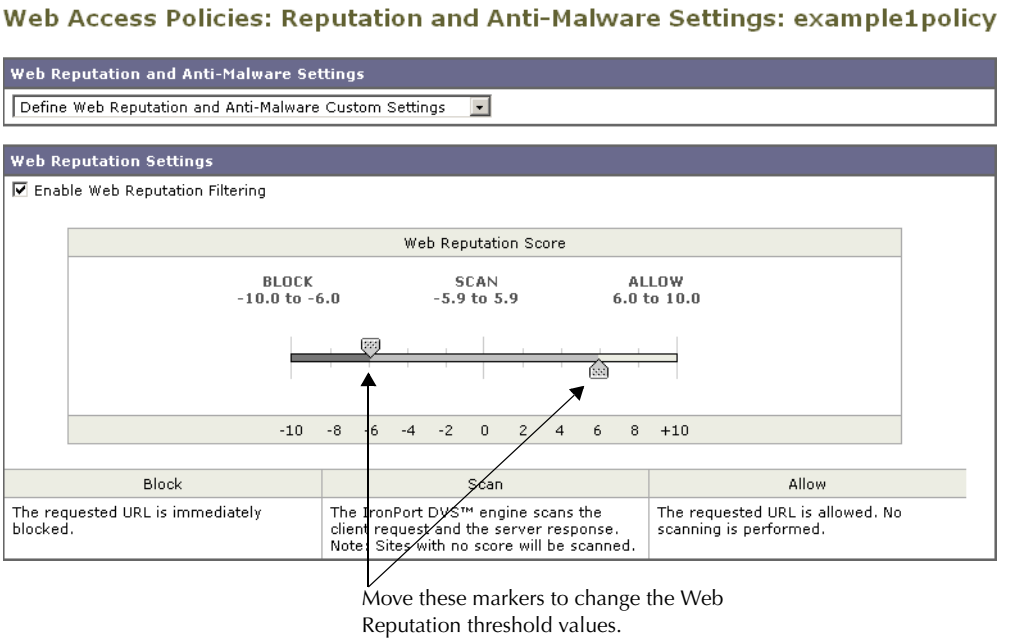

This allows you to override the web reputation and anti-malware settings from the Global Policy Group.

4. Verify the Enable Web Reputation Filtering field is enabled.
5. Move the markers to change the range for URL block, scan, and allow actions.
6. Submit and commit your changes.

Configuring Web Reputation for Decryption Policies

To edit the web reputation filter settings for a decryption policy group:

1. Navigate to the Web Security Manager > Decryption Policies page.
2. Click the link under the Web Reputation column for the decryption policy group you want to edit.
3. Under the Web Reputation Settings section, choose "Define Web Reputation Custom Settings" from the drop down menu if it is not selected already.

   This allows you to override the override the web reputation settings from the Global Policy Group.

Figure 12-2 Web Reputation Filter Settings for Decryption Policies

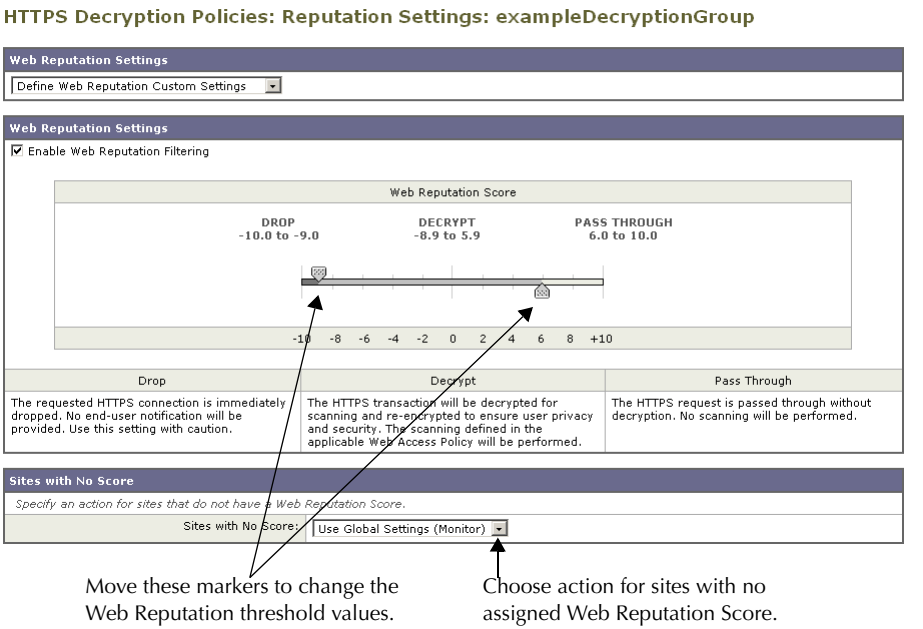

4. Verify the Enable Web Reputation Filtering field is checked.
5. Move the markers to change the range for URL drop, decrypt, and pass through actions.

6. In the Sites with No Score field, choose the action to take on request for sites that have no assigned Web Reputation Score.
7. Submit and commit your changes.

VIEWING WEB REPUTATION FILTERING ACTIVITY

The S-Series appliance supports several options for generating feature specific reports, and displays of summary statistics.

Reports

You can use options on the Monitor > Reports pages of the web interface to select a type of report, capture data, schedule periodic email delivery, and archive reports.

For more information about working with reports, see "Reporting Overview" on page 324.

Monitoring Filter and Scoring Activity

The Monitor > Web Reputation page provides statistical displays of filtering activity. You can update these displays by specifying a time range of hour, day, week or month. Additionally, you have the option to print these display pages and export the raw data to a file.

You can use the following interactive displays and reporting tools to view the results of Web Reputation filtering and scoring activity:

Table 12-3 Web Reputation Filtering Reports

To view...	See...
Web reputation action and scoring activity	Monitor > Web Reputation Filters
Web reputation log files	System Administration > Log Subscriptions • WBRS log files • Access log file

Access Log File

The access log file provides a record of filtering activity. You can examine entries in the access log file and trace web reputation processing specific to individual requests.

For more information about reading access log files, see "Access Log File" on page 343. For more an example access log entry that explains web reputation processing, see "Web Reputation Filters Example" on page 348.

CHAPTER 13

Anti-Malware Services

This chapter contains the following information:

ANTI-MALWARE OVERVIEW

The Web Security appliance anti-malware feature is a security component that uses the IronPort DVS™ engine in combination with the Webroot™ and McAfee technology to identify and stop a broad range of web-based malware threats.

For more information about the DVS engine, see "IronPort DVS™ (Dynamic Vectoring and Streaming) Engine" on page 240.

To use the anti-malware component of the appliance, you must first configure global settings and then apply specific settings to different policies. For more information about configuring the appliance for anti-malware scanning, see "Configuring Anti-Malware Scanning" on page 246.

You can also view the anti-malware scanning activity in reports and in the Web Security Monitor. For more information, see "Viewing Anti-Malware Scanning Activity" on page 250.

Malware Category Descriptions

Table 13-1 describes the different categories of malware the Web Security appliance can block.

Table 13-1 Malware Category Descriptions

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.
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.

Table 13-1 Malware Category Descriptions (Continued)

Malware Type	Description
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.
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.
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.

IRONPORT DVS™ (DYNAMIC VECTORING AND STREAMING) ENGINE

The IronPort Dynamic Vectoring and Streaming (DVS) engine inspects web traffic to provide protection against the widest variety of web-based malware ranging from commercially invasive adware applications, to malicious trojans, system monitors, and phishing attacks.

To configure the DVS engine, and Webroot and McAfee global settings, see "Configuring Anti-Malware Scanning" on page 246.

The IronPort DVS engine can use one or more scanning engines to determine malware risk. Depending on the features purchased with the appliance, you can enable any of the following scanning engines:

- **Webroot.** Webroot's automated spyware detection system rapidly identifies existing and new spyware threats on the Internet by intelligently scanning millions of sites on a daily basis. Webroot uses a signature database to help detect threats on the Internet. For more information about the Webroot scanning engine, see "Webroot Scanning" on page 243.
- **McAfee.** The McAfee scanning engine can detect existing and new malware threats by using a signature database of malware information and heuristic analysis. For more information about the McAfee scanning engine, see "McAfee Scanning" on page 244.

The scanning engines inspect URL transactions to determine a malware scanning verdict to pass to the DVS engine. A malware scanning verdict is a value assigned to a URL request or server response that determines the probability that it contains malware. The DVS engine determines whether to monitor or block the request based on the malware scanning verdicts. For more information about malware scanning verdicts, see "Malware Scanning Verdict Values" on page 353.

In some cases, the DVS engine might determine multiple verdicts for a single URL. For more information about how the DVS handles multiple verdicts, see "Working with Multiple Malware Verdicts" on page 241.

Maintaining the Database Tables

The Webroot and McAfee databases periodically receive updates from the IronPort update server (`https://update-manifests.ironport.com`). Server updates are automated, and the update interval is set by the server, not the appliance. Updates to the database tables occur with a regular degree of frequency, and require no administrator intervention.

For information about update intervals and the IronPort update server, see "Component Updates" on page 391.

How the DVS Engine Works

The DVS engine performs anti-malware scanning on URL transactions that are forwarded from the Web Reputation Filters. Web Reputation Filters calculate the probability that a particular URL contains malware, and assign a URL score that is associated with an action to block, scan, or allow the transaction.

When the assigned web reputation score indicates to scan the transaction, the DVS engine receives the URL request and server response content. The DVS engine, in combination with the Webroot and/or McAfee scanning engines, returns a malware scanning verdict. The DVS engine uses information from the malware scanning verdicts and access policy settings to determine whether to block or deliver the content to the client.

When you enable both Webroot and McAfee, the DVS engine determines how to scan the content to optimize performance and efficacy.

Working with Multiple Malware Verdicts

In some cases, the DVS engine might determine multiple malware verdicts for a single URL. Multiple verdicts can come from one or both scanning engines:

- **Different verdicts from different scanning engines.** When you enable both Webroot and McAfee, each scanning engine might return different malware verdicts for the same object.
- **Different verdicts from the same scanning engine.** A scanning engine might return multiple verdicts for a single object when the object contains multiple infections. For example, a zip file might contain multiple files, each infected with a different kind of malware.

When a URL causes multiple verdicts, the appliance takes different action depending on whether one or both scanning engines return the multiple malware verdicts.

Different Scanning Engines

When a URL causes multiple verdicts from both scanning engines, the appliance performs the most restrictive action. For example, if one scanning engine returns a block verdict and the other a monitor verdict, the DVS engine always blocks the request. Only the most restrictive verdict is logged and reported.

Same Scanning Engine

When a URL causes multiple verdicts from the same scanning engine, the appliance takes action according to the verdict with the highest priority. Only the highest verdict is logged and reported. The following text lists the possible malware scanning verdicts from the highest to the lowest priority.

- Virus
- Trojan Downloader
- Trojan Horse
- Trojan Phisher
- Hijacker
- System monitor
- Commercial System Monitor

- Dialer
- Worm
- Browser Helper Object
- Phishing URL
- Adware
- Encrypted file
- Unscannable
- Other Malware

Suppose the McAfee scanning engine detects both adware and a virus in the scanned object, and that the appliance is configured to block adware and monitor viruses. According to the list above, viruses belong in a higher priority verdict category than adware. Therefore, the appliance *monitors* the object and reports the verdict as virus in the reports and logs. It does not block the object even though it is configured to block adware.

WEBROOT SCANNING

The Webroot scanning engine inspects objects to determine the malware scanning verdict to send to the DVS engine. The Webroot scanning engine inspects the following objects:

- **URL request.** Webroot evaluates a URL request to determine if the URL is a malware suspect. If Webroot suspects the response from this URL might contain malware, the appliance monitors or blocks the request, depending on how the appliance is configured. If Webroot evaluation clears the request, the appliance retrieves the URL and scans the server response.
- **Server response.** When the appliance retrieves a URL, Webroot scans the server response content and compares it to the Webroot signature database.

For more information about how the DVS engine uses malware scanning verdicts to handle web traffic, see "IronPort DVS™ (Dynamic Vectoring and Streaming) Engine" on page 240.

For information about Web Reputation Filtering and URL scores, see "Web Reputation Filters" on page 227.

MCAFEE SCANNING

The McAfee scanning engine inspects objects downloaded from a web server in HTTP responses. After inspecting the object, it passes a malware scanning verdict to the DVS engine so the DVS engine can determine whether to monitor or block the request.

The McAfee scanning engine uses the following methods to determine the malware scanning verdict:

- Matching virus signature patterns
- Heuristic analysis

For more information about how the DVS engine uses malware scanning verdicts to handle web traffic, see "IronPort DVS™ (Dynamic Vectoring and Streaming) Engine" on page 240.

Matching Virus Signature Patterns

McAfee uses virus definitions in its database with the scanning engine to detect particular viruses, types of viruses, or other potentially unwanted software. It searches for virus signatures in files.

When you enable McAfee, the McAfee scanning engine always uses this method to scan server response content.

Heuristic Analysis

New threats on the web appear almost daily. Using only virus signatures, the engine cannot detect a new virus or other malware because its signature is not yet known. However, by using heuristic analysis, the McAfee scanning engine can detect new classes of currently unknown viruses and malware in advance.

Heuristic analysis is a technique that uses general rules, rather than specific rules, to detect new viruses and malware. When the McAfee scanning engine uses heuristic analysis, it looks at the code of an object, applies generic rules, and determines how likely the object is to be virus-like.

Using heuristic analysis increases the likelihood of catching viruses and malware before McAfee updates its virus signature database. However, it also increases the possibility of reporting false positives (clean content designated as a virus). It also might impact appliance performance.

When you enable McAfee, you can choose whether or not to also enable heuristic analysis when scanning objects.

McAfee Categories

Table 13-2 lists the McAfee verdicts and how they correspond to malware scanning verdict categories.

Table 13-2 Appliance Categories for McAfee Verdicts

McAfee Verdict	Malware Scanning Verdict Category
Known Virus	Virus
Trojan	Trojan Horse
Joke File	Adware
Test File	Virus
Wannabe	Virus
Killed	Virus
Commercial Application	Commercial System Monitor
Potentially Unwanted Object	Adware
Potentially Unwanted Software Package	Adware
Encrypted File	Encrypted File

For a list of malware scanning verdicts, see “Malware Scanning Verdict Values” on page 353.

CONFIGURING ANTI-MALWARE SCANNING

The DVS engine and Webroot and McAfee are enabled by default during system setup. Anytime after system setup, you can configure the anti-malware settings for the Web Security appliance. You configure the following anti-malware settings:

- **Global anti-malware settings.** Set object scanning parameters, specify global settings for URL matching, and control when to block the URL or allow processing to continue.
- **Access policy anti-malware settings.** Enable monitoring or blocking for malware categories based on malware scanning verdicts.

To configure anti-malware settings:

1. On the Security Services > Anti-Malware page, click **Edit Global Settings**.

   The Edit Anti-Malware Settings page appears.

2. Configure the anti-malware settings as necessary. Table 13-3 describes the anti-malware settings you can configure.

Table 13-3 Anti-Malware Settings

Setting	Description
Object Scanning Limits	Specify a maximum object size and timeout value for single objects.
Domain Levels for Malware Request Detection	This value specifies the number of domain name elements to match when processing a URL. If the URL matches a hostname in the Webroot signature database, URL checking continues to match the number of domain name elements specified in this parameter. Valid range for this parameter is 3-100 where a minimum value of 8 is recommended to avoid a level of matching that results in inaccurately blocked web sites. Applies to the Webroot scanning engine only.

Table 13-3 Anti-Malware Settings (Continued)

Setting	Description
Threat Risk Threshold	The TRT (Threat Risk Threshold) assigns a numerical value to the probability that malware exists. Proprietary algorithms evaluate the result of a URL matching sequence and assign a TRR (Threat Risk Rating). This value is associated with the threat risk threshold setting. If the TRR value is greater than or equal to the TRT, the URL is considered malware and is passed on for further processing. Note: Setting the Threat Risk Threshold to a value lower than 90 dramatically increases the rate of URL blocking and denies legitimate requests. IronPort strongly recommends maintaining the TRT default value of 90. The minimum value for a TRT setting is 51. Applies to the Webroot scanning engine only.
Heuristic Scanning	Choose whether or not to enable heuristic scanning for the McAfee scanning engine. For more information about heuristic scanning, see "McAfee Scanning" on page 244. **Note:** Heuristic analysis increases security protection, but can result in false positives and decreased performance. Applies to the McAfee scanning engine only.

3. Submit and commit your changes.
4. Navigate to the Web Security Manager > Access Policies page.
5. Click the Web Reputation and Anti-Malware Filtering link for the access policy you want to configure.

   On this page, you can enable monitoring or blocking for malware categories based on malware scanning verdicts.
6. Under the "Web Reputation and Anti-Malware Settings" section, choose Define Web Reputation and Anti-Malware Custom Settings if it is not chosen already.

   This allows you to configure web reputation and anti-malware settings for this access policy that differ from the global policy.
7. Scroll down to the Ironport DVS Anti-Malware Settings section.

Figure 13-1 Access Policy Anti-Malware Settings

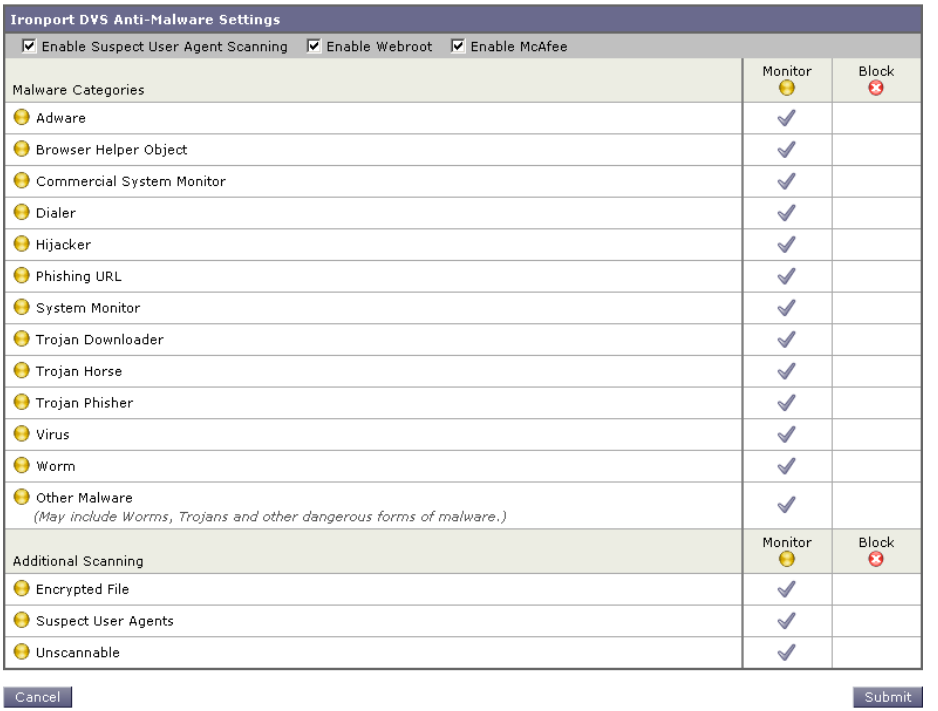

8. Configure the anti-malware settings for the policy as necessary. Table 13-4 describes the anti-malware settings you can configure for access policies.

Table 13-4 Anti-Malware Settings for Access Policies

Setting	Description
Enable Suspect User Agent Scanning	Choose whether or not to enable the appliance to scan traffic based on the user agent field specified in the HTTP request header. When you check this setting, you can choose to monitor or block suspect user agents in the Additional Scanning section at the bottom of the page.
Enable Webroot	Choose whether or not to enable the appliance to use the Webroot scanning engine when scanning traffic. When you enable Webroot scanning, you can choose to monitor or block some additional categories in the Malware categories on this page.

Table 13-4 Anti-Malware Settings for Access Policies (Continued)

Setting	Description
Enable McAfee	Choose whether or not to enable the appliance to use the McAfee scanning engine when scanning traffic. When you enable McAfee scanning, you can choose to monitor or block some additional categories in the Malware categories on this page.
Malware Categories	Choose whether to monitor or block the various malware categories based on a malware scanning verdict. The categories listed in this section depend on which scanning engines you enable above.
Additional Scanning	Choose whether to monitor or block the types of objects and responses listed in this section. **Note:** URL transactions are categorized as unscannable when the configured maximum time setting is reached or when the system experiences a transient error condition. For example, transactions might be categorized as unscannable during scanning engine updates or AsyncOS upgrades. The malware scanning verdicts SV_TIMEOUT and SV_ERROR, are considered unscannable transactions.

9. Submit and commit your changes.

VIEWING ANTI-MALWARE SCANNING ACTIVITY

The Web Security appliance supports several options for generating feature specific reports, and interactive displays of summary statistics.

Reports

You can use options on the Monitor > Reports pages of the web interface to select a type of report, capture data, schedule periodic email delivery, and archive reports.

For more information about working with reports, see "Reporting Overview" on page 324.

Monitoring Scanning Activity

The Monitor > Anti-Malware page provides statistical displays of malware scanning activity. You can update these displays by specifying a time range of hour, day, week or month. Additionally, you have the option to print these display pages and export the raw data to a file. You can use the following interactive displays and reporting tools to view the results of anti-malware scanning and related activity:

Table 13-5 Anti-Malware Scanning Reports

To View...	See...
Top anti-malware sites	Monitor > Overview
Top malware categories and threats	Monitor > Anti-Malware
Anti-malware log files	System Administration > Log Subscriptions • Webroot log files • McAfee log files • Access log file

Access Log File

The access log file provides a record of anti-malware scanning activity. You can examine entries in the access log file and trace the result of malware scanning for individual requests. For more information about reading access log files, see "Access Log File" on page 343.

CHAPTER 14

Authentication

This chapter contains the following information:

AUTHENTICATION OVERVIEW

Authentication is the act of confirming the identity of a user. By using authentication in the Web Security appliance, you can control access to the Web for each user or a group of users. This allows you to enforce the organization's policies and comply with regulations. When you enable authentication, the Web Security appliance authenticates clients on the network before allowing them to connect to a destination server.

The Web Security appliance supports the following authentication protocols:

- **Lightweight Directory Access Protocol (LDAP).** The appliance supports standard LDAP server authentication and secure LDAP authentication. You can use a Basic authentication scheme. For more information about LDAP configuration options, see "Authenticating Using LDAP" on page 281.
- **NT Lan Manager (NTLM).** The appliance supports NTLM to enable authentication between the appliance and a Microsoft Windows domain controller. You can use either NTLMSSP or Basic authentication schemes. For more information about NTLM configuration options, see "Authenticating Using NTLM" on page 286.

To enable authentication, you must create at least one authentication realm. An authentication realm is a set of authentication servers (or a single server) supporting a single authentication protocol with a particular configuration. For more information about authentication realms, see "Working with Authentication Realms" on page 262.

When you create more than one realm, you can group the realms into an authentication sequence. An authentication sequence is a group of authentication realms listed in the order the Web Security appliance uses for authenticating clients. For more information about authentication sequences, see "Working with Authentication Sequences" on page 264.

You configure some authentication options at a global level, independent of any realm. For more information, see "Configuring Global Authentication Settings" on page 271.

By creating authentication realms and sequences, you can configure the Web Security appliance to use one or more authentication servers for authenticating clients on the network. For more information about how the appliance works when it uses multiple authentication servers, see "Appliance Behavior with Multiple Authentication Realms" on page 267.

After creating an authentication realm and possibly a sequence, too, you can create or edit identities based on authentication realms or sequences. Note, however, that if you delete an authentication realm or sequence, any identity policy group that depends on the deleted realm or sequence becomes disabled. For more information about using authentication with identities, see "How Authentication Affects Identity Groups" on page 106.

Client Application Support

When the Web Security appliance is deployed in transparent mode and a transaction requires authentication, the Web Proxy replies to the client application asking for authentication credentials. However, not all client applications support authentication, so they have no

method for prompting users to provide their user names and passwords. These applications cannot be used when the Web Security appliance is deployed in transparent mode.

The following is a partial list of applications that do not work when the appliance is deployed in transparent mode:

- Mozilla Thunderbird
- Adobe Acrobat Updates
- HttpBridge
- Subversion, by CollabNet
- Microsoft Windows Update
- Microsoft Visual Studio

Note — If users need to access a particular URL using one of these client applications, then create an identity based on a custom URL category that does not require authentication and place the identity above all other identities that require authentication. When you do this, the client application will not be asked for authentication.

Authenticating Users

When users access the web through the Web Security appliance, they might get prompted to enter a user name and password. The Web Proxy requires authentication credentials for some users depending on the configured identity and access policy groups. Users should enter the user name and password of the credentials recognized by the organization's authentication server.

When the Web Proxy uses NTLMSSP authentication with an NTLM authentication realm, users are typically not prompted to enter a user name and password if single sign-on is configured correctly. However, if users are prompted for authentication, they must type the name of their Windows domain before their user name. For example, if user jsmith is on Windows domain MyDomain, then the user should type the following text in the user name field:

```
MyDomain\jsmith
```

However, if the Web Proxy uses Basic authentication for an NTLM authentication realm, then entering the Windows domain is optional. If the user does not enter the Windows domain, then the Web Proxy will prepend the default Windows domain automatically.

Note — When the Web Proxy uses authentication with an LDAP authentication realm, users should not enter the Windows domain name.

Working with Upstream Proxy Servers

When the Web Security appliance is connected to an upstream proxy server, you can configure the appliance or the upstream proxy to use authentication, but not both. IronPort

recommends configuring the Web Security appliance to use authentication. This allows you to create policies based on user authentication.

If both the appliance and the upstream proxy use authentication, depending on the configurations, the appliance and upstream proxy might engage in an infinite loop of requesting authentication credentials. For example, if the upstream proxy requires Basic authentication, but the appliance requires NTLMSSP authentication, then the appliance can never successfully pass Basic credentials to the upstream proxy. This is due to limitations in authentication protocols.

HOW AUTHENTICATION WORKS

To authenticate users accessing the Web, the Web Security appliance connects to an external authentication server. The authentication server contains a list of users and their corresponding passwords and groups the users into a hierarchy. For users on the network to successfully authenticate, they must provide valid authentication credentials (user name and password as stored in the authentication server).

When users access the Web through a Web Security appliance requiring authentication, the Web Proxy asks the client for authentication credentials. The Web Proxy communicates with both the client and the authentication server to authenticate the user and eventually process the original request.

Figure 14-1 shows how the Web Security appliance communicates with clients and authentication servers.

Figure 14-1 Web Security Appliance Authentication

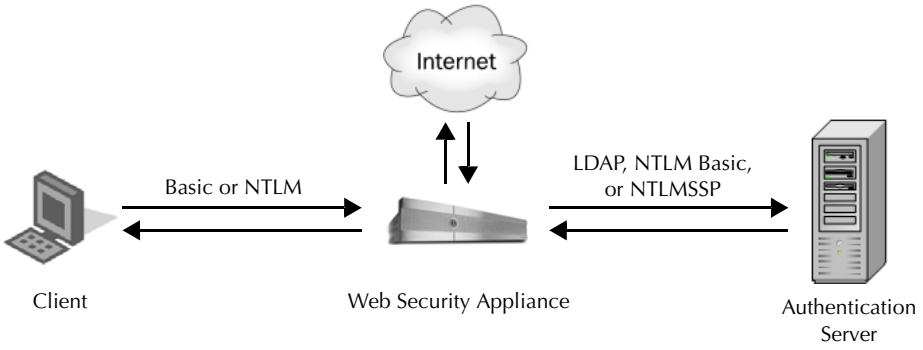

The Web Security appliance supports the following authentication protocols:

- **Lightweight Directory Access Protocol (LDAP).** The Web Proxy uses the LDAP Bind operation to query an LDAP compatible authentication server. The appliance supports standard LDAP server authentication and secure LDAP authentication.

  For more information about LDAP configuration options, see “Authenticating Using LDAP” on page 281.

- **NT Lan Manager (NTLM).** The NTLM protocol works with Microsoft Active Directory servers. It uses a challenge-response sequence requiring the transmission of multiple messages between the client and the Active Directory server.

  For more information about NTLM configuration options, see “Authenticating Using NTLM” on page 286.

In addition to the supported protocols listed above, the Web Security appliance supports the following authentication schemes:

- **Basic.** Allows a client application to provide authentication credentials when making a request. You can use the Basic authentication scheme with either an LDAP or Active Directory server.
- **NTLMSSP.** Uses a binary message format to authenticate clients that use the NTLM protocol to access network resources. You can use the NTLMSSP authentication scheme with an Active Directory server. When the Web Proxy uses NTLMSSP, most client applications can authenticate automatically using the Windows login credentials without prompting users for their credentials again. This is called "single sign-on."

For more information, see "Basic versus NTLMSSP Authentication Schemes" on page 256.

Table 14-1 describes the different authentication scenarios you can configure between the Web Security appliance and the client and between the Web Security appliance and the authentication server.

Table 14-1 Web Security Appliance Authentication Scenarios

Client to Web Security Appliance	Web Security Appliance to Authentication Server	Authentication Server Type
Basic	LDAP	LDAP server
Basic	LDAP	Active Directory server using LDAP
Basic	NTLM Basic	Active Directory server (NTLM Basic)
NTLM	NTLMSSP	Active Directory server (NTLMSSP)

Web Proxy deployment also affects how authentication works for the scenarios described in Table 14-1. For more information, see "How Web Proxy Deployment Affects Authentication" on page 257.

Basic versus NTLMSSP Authentication Schemes

When you configure an identity group to use authentication, you can choose the authentication scheme to use, either Basic or NTLMSSP. The authentication scheme used affects what the user experiences and the security of the password entered by the user.

Table 14-2 describes the differences between Basic and NTLMSSP authentication schemes.

Table 14-2 Basic versus NTLMSSP Authentication Schemes

Authentication Scheme	User Experience	Security
Basic	The client always prompts users for credentials. After the user enters credentials, browsers typically offer a check box to remember the provided credentials. Any time the browser is closed, the client either prompts again or resends the previously remembered credentials again.	Credentials are sent *insecurely* using clear text (Base64). A simple packet capture between the client and the Web Security appliance can reveal the user name and password.
NTLMSSP	The client transparently authenticates using its Windows logon credentials, meaning the user does not get prompted for credentials. However, the client prompts the user for credentials under the following circumstances: • The Windows credentials failed. • The client does not trust the Web Security appliance.	Credentials are sent *securely* using a 3-way handshake (digest style authentication). The password is never sent across the connection. For more information on the 3-way handshake, see "Explicit Forward Deployment, NTLM Authentication" on page 259.

How Web Proxy Deployment Affects Authentication

The Web Proxy communicates with clients and authentication servers differently depending on how it is deployed and the authentication protocol configured.

Table 14-3 lists the possible authentication configurations when considering authentication protocol and deployment method.

Table 14-3 Methods of Authentication

Web Proxy Deployment	Client to Web Security Appliance	Web Security Appliance to Authentication Server
Explicit forward	Basic	LDAP or NTLM Basic
Transparent	Basic	LDAP or NTLM Basic
Explicit forward	NTLM	NTLMSSP
Transparent	NTLM	NTLMSSP

This section describes each row in Table 14-3 in more detail.

Explicit Forward Deployment, Basic Authentication

When a client explicitly sends a web page request to a Web Security appliance deployed in explicit forward mode, the Web Proxy can reply to the client with a 407 HTTP response "Proxy Authentication Required." This status informs the client that to access web resources, it must supply valid authentication credentials.

1. Client sends a request to the Web Proxy to connect to a web page.
2. Web Proxy responds with a 407 HTTP response "Proxy Authentication Required."
3. User enters credentials, and the client application resends the original request with the credentials encoded in Base64 (not encrypted) in a "Proxy-Authorization" HTTP header.
4. The Web Proxy verifies the credentials and returns the requested web page.

Table 14-4 lists some advantages and disadvantages of using explicit forward Basic authentication.

Table 14-4 Advantages and Disadvantages of Explicit Forward Basic Authentication

Advantages	Disadvantages
• RFC Based • Supported by all browsers, most other programs • Minimal overhead • Works for HTTPS (CONNECT) requests	• Password passed in clear text (Base64) on every request • No single sign-on

Transparent Deployment, Basic Authentication

407 HTTP responses are only allowed from proxy servers. However, when the Web Proxy is deployed in transparent mode, client applications on the network are not made aware of the existence of a proxy. Therefore, the Web Proxy cannot return a 407 response.

To get around this:

1. Client sends a request to a web page, but the Web Proxy transparently intercepts it.
2. The Web Proxy redirects the client to the Web Proxy pretending to be a local web server using a 307 HTTP response.
3. The client sends a request to the redirected URL.
4. The Web Proxy sends a 401 HTTP response "Authorization required."
5. The user gets prompted for credentials and enters them.
6. The client sends the request again, but this time with the credentials in a "Authorization" HTTP header.
7. The Web Proxy confirms the credentials, keeps track of the user using either an IP address or a cookie, and then redirects the client to the originally requested server.

Note — You can configure whether the Web Proxy uses the IP address or a cookie to track authenticated users.

8. When the client requests the original web page again, the Web Proxy transparently intercepts the request, confirms the user by IP address or cookie, and returns the requested page.

Note — If the client tries to connect to another web page and the Web Proxy tracked the user by IP address, the Web Proxy confirms the user by IP address and returns the requested page.

Table 14-5 lists some advantages and disadvantages of using transparent Basic authentication with IP based credential caching.

Table 14-5 Advantages and Disadvantages of Transparent Basic Authentication—IP Caching

Advantages	Disadvantages
• Works with all major browsers • To use with user agents that do not support authentication at all, users only need to authenticate first in a supported browser • Relatively low overhead • Works for HTTPS requests as long as the user has already authenticated with an HTTP request	• Authentication credentials are tied with the IP address, not the user (does not work in Citrix/ RDP environments, or if the user changes IP address) • No single sign-on • Password is sent in clear text (Base64)

Table 14-15 lists some advantages and disadvantages of using transparent Basic authentication with cookie based credential caching.

Table 14-6 Advantages and Disadvantages of Transparent Basic Authentication—Cookie Caching

Advantages	Disadvantages
• Works with all major browsers • Authentication stays with the browser (that is, the user) rather than the host or IP address	• High overhead: each new web page goes through the entire process • Requires cookies to be enabled • Does not work for HTTPS requests • No single sign-on • Password is sent in clear text (Base64)

Explicit Forward Deployment, NTLM Authentication

The Web Proxy uses a third party challenge and response system to authenticate users on the network.

1. Client sends a request to the Web Proxy to connect to a web page.

1. Web Proxy responds with a 407 HTTP response “Proxy Authentication Required”.
2. Clients repeats request and includes a “Proxy-Authorization” HTTP header with an NTLM “negotiate” message.
3. Web Proxy responds with a 407 HTTP response and an NTLM “challenge” message based on the negotiate message from the client.
4. Client repeats the request and includes a response to the challenge message.

   Note — The client uses an algorithm based on its password to modify the challenge and sends the challenge response to the Web Proxy.

5. Web Proxy passes the authentication information to the Active Directory server. The Active Directory server then verifies that the client is using the correct password based on whether or not it modified the challenge string appropriately.
6. Presuming challenge response passes, the Web Proxy returns the requested web page.

Note — Further requests on the *same TCP connection* do not need to be authenticated again with the Active Directory server.

Table 14-7 lists some advantages and disadvantages of using explicit forward NTLM authentication.

Table 14-7 Advantages and Disadvantages of Explicit Forward NTLM Authentication

Advantages	Disadvantages
• Password is not transmitted to the authentication server, making it more secure • Connection is authenticated, not the host or IP address • True single sign-on in an Active Directory environment when the client applications are configured to trust the Web Security appliance	• Medium overhead: each new connection needs to be re-authenticated • Primarily only supported on Windows and with major browsers

Transparent Deployment, NTLM Authentication

Transparent NTLM authentication is similar to transparent Basic authentication except that the Web Proxy communicates with clients using NTLM instead of Basic. However, with transparent NTLM authentication, the authentication credentials are not sent in the clear to the authentication server.

For more information, see “Transparent Deployment, Basic Authentication” on page 258.

The advantages and disadvantages of using transparent NTLM authentication are the same as those of using transparent Basic authentication except that transparent NTLM authentication is better because the password is not sent to the authentication server and you can achieve

single sign-on when the client applications are configured to trust the Web Security appliance. For more information on the advantages and disadvantages of transparent Basic authentication, see Table 14-5 on page 259 Table 14-6 on page 259.

WORKING WITH AUTHENTICATION REALMS

An authentication realm is a set of authentication servers (or a single server) supporting a single authentication protocol with a particular configuration.

You can perform any of the following tasks when configuring authentication:

- Include up to three authentication servers in a realm.
- Create zero or more LDAP realms.
- Create zero or one NTLM realm.
- Include an authentication server in multiple realms.
- Include one or more realms in an authentication sequence.
- Include realms of different protocols in a single authentication sequence.
- Assign a realm or a sequence to an access policy group.

You create, edit, and delete authentication realms on the Network > Authentication page under the Authentication Realms section. Figure 14-2 shows where you define authentication realms.

Figure 14-2 Authentication Page — Authentication Realms

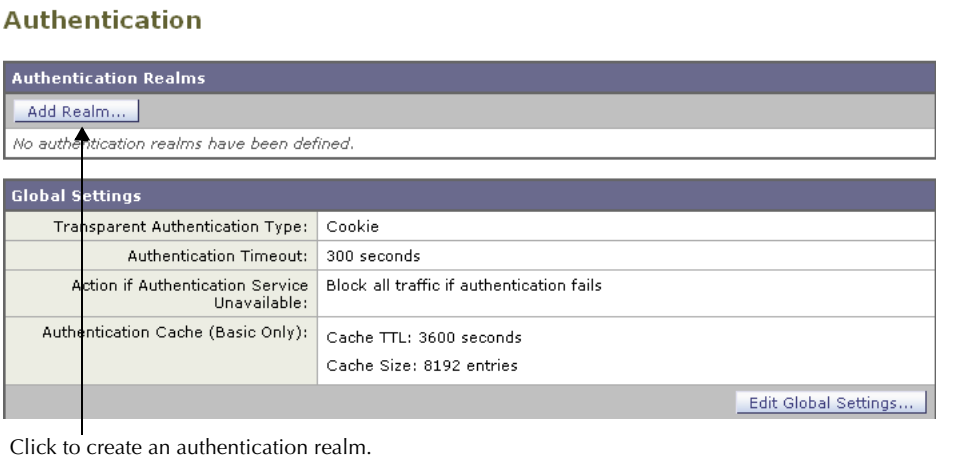

When you create two or more realms, you can order them in an authentication sequence. For more information, see “Working with Authentication Sequences” on page 264.

Creating Authentication Realms

When you first create a realm, you choose the protocol type, either LDAP or NTLM. After you create an NTLM realm, the appliance only allows you to create LDAP realms. After you enter the authentication settings, you can verify that the parameters you entered are valid before

you submit your changes. For more information about testing the authentication settings, see "Testing Authentication Settings" on page 268.

To create an authentication realm:

1. On the Network > Authentication page, click **Add Realm**. The Add Realm page appears.
2. Enter a name for the authentication realm in the Realm Name field.

   Note — All sequence and realm names must be unique.

3. If no NTLM realm is defined, choose the authentication protocol and scheme in the Authentication Protocol and Scheme(s) field.
4. Enter the authentication settings as necessary, depending on the protocol type.
   - For details on LDAP settings, see Table 14-11 on page 281.
   - For details on NTLM settings, see Table 14-12 on page 287.
5. You can test the parameters you entered by clicking **Start Test** in the Test Current Settings section.
6. Submit and commit your changes.

Editing Authentication Realms

To edit an authentication realm:

1. On the Network > Authentication page, click the realm name.
2. Change the name of the realm if necessary.
3. Edit the authentication settings as necessary, depending on the protocol type.
   - For details on LDAP settings, see Table 14-11 on page 281.
   - For details on NTLM settings, see Table 14-12 on page 287.
4. You can test the parameters you entered by clicking **Start Test** in the Test Current Settings section.
5. Submit and commit your changes.

Deleting Authentication Realms

When you delete a realm, the Web Security appliance automatically deletes that realm from any sequence that used it. Also, any access policy group that depends on the deleted realm becomes disabled.

To delete an authentication realm:

1. On the Network > Authentication page, click the trash can icon for the realm name.
2. Confirm that you want to delete the realm by clicking **Delete**.
3. **Commit** your changes.

WORKING WITH AUTHENTICATION SEQUENCES

When you create more than one realm, you can group the realms into an authentication sequence. An authentication sequence is a group of authentication realms listed in the order the Web Security appliance uses for authenticating clients.

You can perform any of the following tasks when configuring authentication sequences:

- Create multiple authentication sequences.
- Include one or more realms in an authentication sequence.
- Include realms of different protocols in a single authentication sequence.
- Assign a realm or a sequence to an access policy group.

You create authentication sequences on the Network > Authentication page under the Realm Sequences section. the Realm Sequences section only appears when you create two or more realms. Figure 14-3 shows where you create, edit, and delete authentication sequences Figure 14-3.

Figure 14-3 Authentication Page — Authentication Sequences

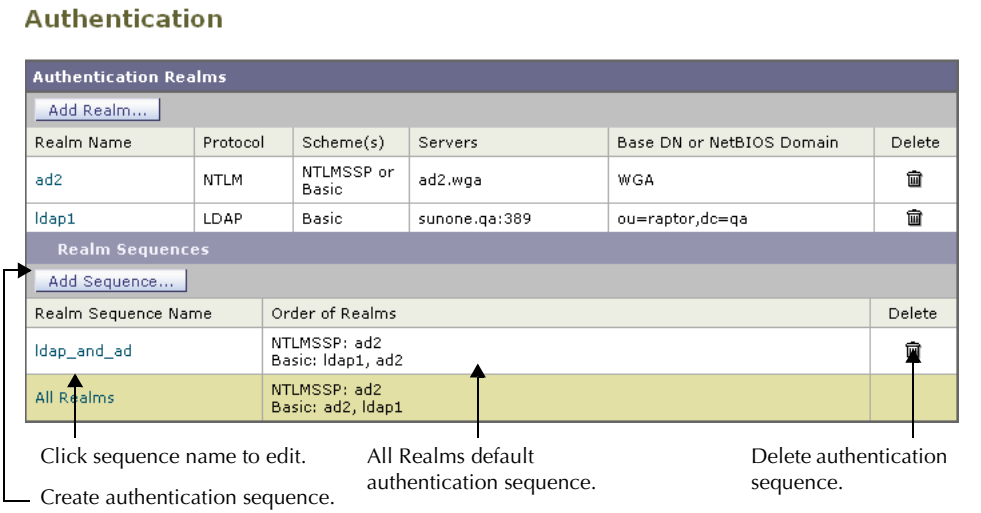

After you create the second realm, the appliance automatically displays the Realm Sequences section and includes a default authentication sequence named All Realms. The All Realms sequence automatically includes each realm you define. You can change the order of the realms within the All Realms sequence, but you cannot delete any of its realms. You cannot delete the All Realms sequence.

Creating Authentication Sequences

You can create an authentication sequence after you create multiple authentication realms.

To create an authentication sequence:

1. On the Network > Authentication page, click **Add Sequence**.

   The Add Realm Sequence page appears.

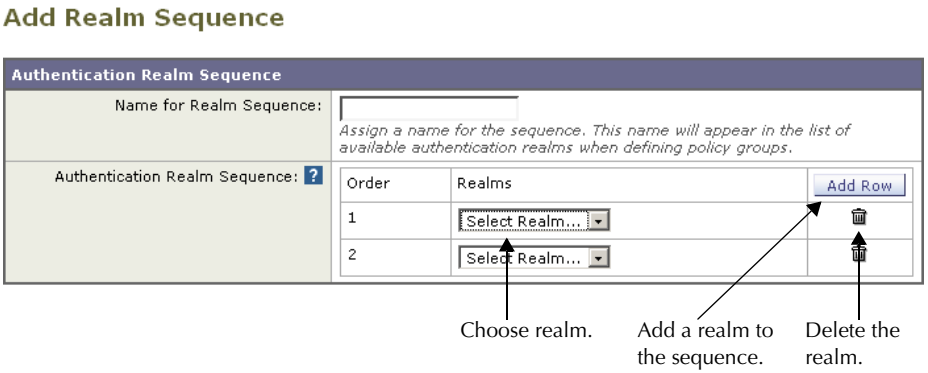

2. Enter a name for the sequence in the Name for Realm Sequence field.

   Note — All sequence and realm names must be unique.

3. In the first row of the Authentication Realm Sequence area, choose the realm name you want to include in the sequence from the Realms field.
4. If you want to include more realms, click **Add Row**.
5. Choose the realm name for any additional row you add.

   Note — You can delete a realm from the sequence by clicking the trash can icon for that row.

6. When you have entered all realms in the sequence, and they are in the order you want, submit and commit your changes.

Editing Authentication Sequences

To edit an authentication sequence:

1. On the Network > Authentication page, click the sequence name.
2. Perform any of the following tasks as necessary:
   - Change the name of the sequence.
   - Add a new realm by clicking **Add Row**.

- Delete a realm by clicking the trash can icon.
- Change the order of the realms by clicking the arrow icon in the Order column for the realm.

3. Submit and commit your changes.

Deleting Authentication Sequences

If you delete an authentication sequence, any access policy group that depends on the deleted sequence becomes disabled.

To delete an authentication sequence:

1. On the Network > Authentication page, click the trash can icon for the sequence name.
2. Confirm that you want to delete the sequence by clicking **Delete**.
3. **Commit** your changes.

APPLIANCE BEHAVIOR WITH MULTIPLE AUTHENTICATION REALMS

You can configure the Web Security appliance to attempt authenticating clients against multiple authentication servers, and against authentication servers with different authentication protocols. When you configure the appliance to authenticate against multiple authentication servers, it only requests the credentials from the clients once. This is true even when you configure the appliance to authenticate against different protocols.

You might want to configure a web policy group to authenticate against different realms if your organization acquires another organization that has its own authentication server using the same or a different authentication protocol. That way, you can create one access policy group for all users and assign to the policy group an authentication sequence that contains a realm for each authentication server.

When you assign an authentication sequence with multiple realms to a policy group and a client sends a content request, the appliance performs the following actions:

1. The appliance gets the credentials from the client.
2. The appliance attempts to authenticate the client against the authentication server(s) defined in the first realm in the sequence.
3. If the client credentials do not match a user in the servers defined in the first realm, it tries to authenticate against the authentication server(s) in the next realm in the sequence.
4. The appliance continues trying to authenticate the client against servers in the next realms until it either succeeds or runs out of authentication realms.
5. When authentication succeeds, the appliance passes the content response to the client.
6. When the appliance fails to authenticate the client against any authentication realm in the sequence, the appliance does not allow the client to connect to the destination server. Instead, it displays an error message to the client.

Tip: For optimal performance, configure clients on a subnet to be authenticated in a single realm.

TESTING AUTHENTICATION SETTINGS

When you create or edit an authentication realm, you enter a lot of configuration settings to connect to the authentication server. You can test the settings you enter before submitting the changes to verify you entered the connection information correctly.

You can test authentication setting from either the CLI or the web interface:

- **Web interface.** Use **Start Test** when you create or edit an authentication realm. For more information, see "Testing Authentication Settings in the Web Interface" on page 269.
- **CLI command.** Use the `testauthconfig` command. For more information, see "Testing Authentication Settings in the CLI" on page 270.

Testing Process

When you test authentication settings, the Web Security appliance first verifies that the settings you entered for the realm are in valid formats. For example, if a field requires a string and it currently contains a numeric value, the appliance informs you of that error.

If all fields contain valid values, the appliance performs different steps, depending on the authentication protocol. If the realm contains multiple authentication servers, the appliance goes through the testing process for each server in turn.

The appliance continues testing all servers in the realm and determines as many failures as possible for each server. It reports the testing outcome of each server in the realm.

LDAP Testing

The appliance performs the following steps when you test LDAP authentication settings:

1. It ensures that the LDAP server is listening on the specified LDAP port.
2. If Secure LDAP is selected, the appliance ensures the LDAP server supports secure LDAP.
3. It performs an LDAP query using the supplied Base DN, User Name Attribute, and User Filter Query.
4. If the realm includes Bind Parameters, the appliance validates them by forming an LDAP query with the Bind Parameters.
5. If Group Authorization is provided, the appliance ensures that the specified group attributes are valid by fetching the groups from the server.

NTLM Testing

The appliance performs the following steps when you test NTLM authentication settings:

1. It ensures that the specified Active Directory server is reachable and responds to queries.
2. It ensures that a DNS lookup on the Active Directory domain is successful since it must be a DNS domain name and not a WINS domain name.
3. It validates the user credentials by generating a kerberos ticket.

4. It validates whether the user has the proper privileges to add the Web Security appliance to the Active Directory domain.
5. It validates whether you can fetch the groups within the domain.

Testing Authentication Settings in the Web Interface

You verify the authentication settings in the Test Current Settings section when you create or edit an authentication realm.

Figure 14-4shows where you verify the authentication settings in the web interface.

Figure 14-4 Network > Authentication Page — Test Current Settings Section

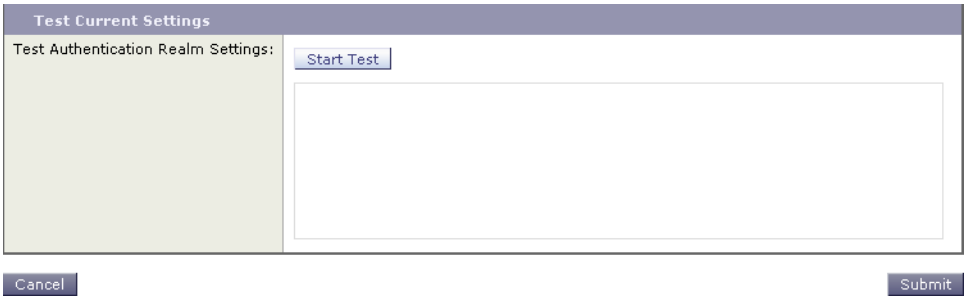

After you enter all settings, click **Start Test**. The appliance uses the connection information entered to attempt to connect to the authentication server. It displays the status of the test below **Start Test**.

Start Test changes to **Stop Test** while the appliance tests the settings against the authentication servers. If the testing takes too much time and you already know it is going to fail, you can click **Stop Test** to stop the testing process and edit the settings.

Figure 14-5 shows the testing results for an LDAP authentication realm.

Figure 14-5 Authentication Testing Results

Testing Authentication Settings in the CLI

You can use the `testauthconfig` CLI command to test authentication settings defined for a given realm. The command syntax is:

```
testauthconfig [-d level] [realm name]
```

Running the command without any option causes the appliance to list the configured authentication realms from which you can make a selection.

The debug flag (`-d`) controls the level of debug information. The levels can range between 0-10. If unspecified, the appliance uses a level of 0. With level 0, the command will return success or failure. If the test settings fail, the command will list the cause of the failure.

Note — IronPort recommends you use level 0. Only use a different debug level when you need more detailed information to troubleshoot.

For more information about the `testauthconfig` command, see "Web Security Appliance CLI Commands" on page 422.

CONFIGURING GLOBAL AUTHENTICATION SETTINGS

Some authentication settings are independent of any realm you define. For example, you can configure whether or not clients send authentication credentials to the Web Security appliance securely, even when using Basic authentication scheme. For more information, see "Sending Authentication Credentials Securely" on page 279.

Figure 14-6 shows the global authentication settings on the Network > Authentication page.

Figure 14-6 Authentication Global Settings

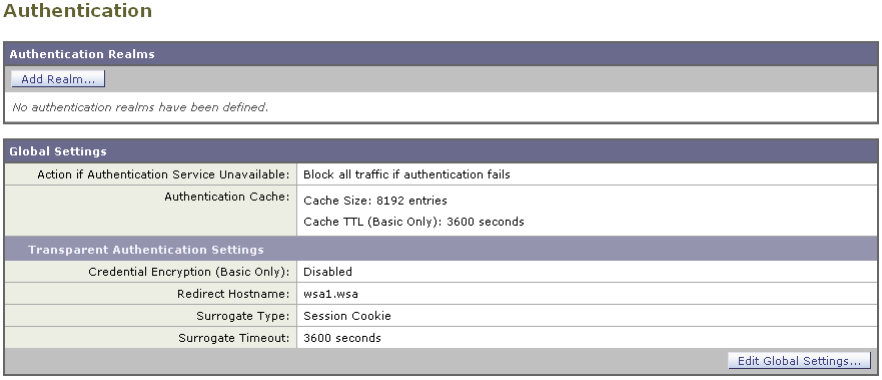

Note — The global authentication settings you can configure changes according to the Web Proxy deployment. You can configure more settings when it is deployed in transparent mode than in explicit forward mode.

To configure global authentication settings:

1. On the Network > Authentication page, click **Edit Global Settings**.

   Figure 14-7 on page 272 shows the Edit Global Authentication Settings page.

Figure 14-7 Global Authentication Settings

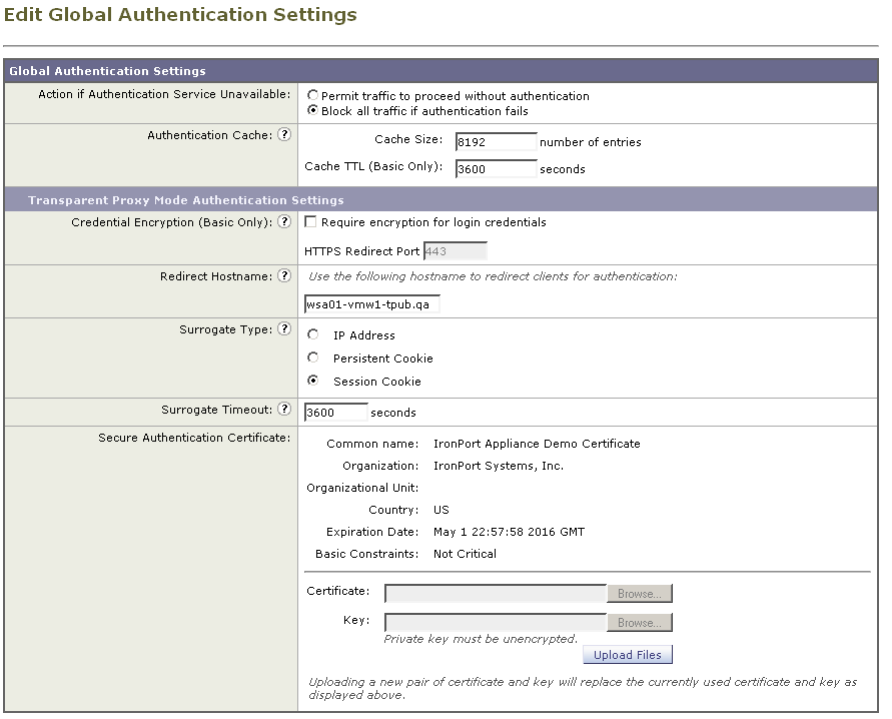

2. Edit the settings defined in Table 14-8.

Table 14-8 Global Authentication Settings

Setting	Description
Action if Authentication Service Unavailable	Choose one of the following values: • **Permit traffic to proceed without authentication.** Processing continues as if the user was authenticated. • **Block all traffic if user authentication fails.** Processing is discontinued and all traffic is blocked.

Table 14-8 Global Authentication Settings (Continued)

Setting	Description
Authentication Cache	These settings determine how long the Web Security appliance caches the client user name and password before revalidating them with the authentication server. Configure the following options: • **Cache Size.** Specifies the number of entries that are stored in the authentication cache. Set this value to safely accommodate the number of users that are actually using this device. The default value is the recommended setting. • **Cache TTL (Basic only).** Controls the length of time that user credentials are stored in the cache. The default value is the recommended setting. When the Surrogate Timeout setting is configured and is greater than the Cache TTL, then the Surrogate Timeout value takes precedence and the Web Proxy contacts the authentication server after surrogate timeout expires. **Note:** In explicit forward mode, you can configure the surrogate timeout when you enable secure client authentication or from the `advancedproxyconfig > authentication` CLI command.

The remaining authentication settings you can configure depends on how the Web Proxy is deployed, in transparent or explicit forward mode.

Figure 14-8 on page 274 shows where you configure the global authentication settings when the Web Proxy is deployed in transparent mode.

Figure 14-8 Transparent Proxy Mode Authentication Settings

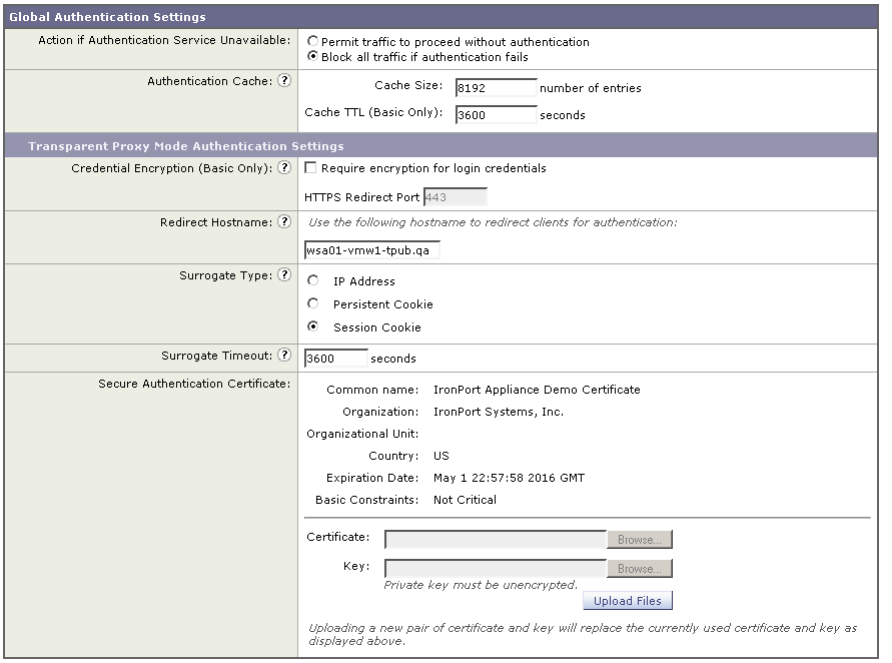

3. If the Web Proxy is deployed in transparent mode, edit the settings in Table 14-9.

Table 14-9 Transparent Proxy Mode Authentication Settings

Setting	Description
Credential Encryption (Basic Only)	Choose whether or not the Web Proxy redirects clients to securely pass authentication credentials to the Web Proxy using HTTPS. When you enable this setting, you must also specify which port to use for redirecting requests using HTTPS. This setting applies to requests using the Basic authentication scheme only since NTLMSSP requests are secure by definition. For more information, see “Sending Authentication Credentials Securely” on page 279.

Table 14-9 Transparent Proxy Mode Authentication Settings (Continued)

Setting	Description
Redirect Hostname	Enter the short host name of the network interface on which the Web Proxy listens for incoming connections. When you configure authentication on an appliance deployed in transparent mode, the Web Proxy uses this host name in the redirection URL sent to clients for authenticating users. You can enter either the following values: • **Single word host name.** You can enter the single word host name that is DNS resolvable by the client and the Web Security appliance. This allows clients to achieve true single sign-on with Internet Explorer without additional browser side setup. Be sure to enter the single word host name that is DNS resolvable by the client and the Web Security appliance. For example, if your clients are in domain `mycompany.com` and the interface on which the Web Proxy is listening has a full host name of `proxy.mycompany.com`, then you should enter `proxy` in this field. Clients perform a lookup on `proxy` and they should be able to resolve `proxy.mycompany.com`. • **Fully qualified domain name (FQDN).** You can also enter the FQDN or IP address in this field. However, if you do that and want true single sign-on for Internet Explorer and Firefox browsers, you must ensure that the FQDN or IP address is added to the client's Trusted Sites list in the client browsers. The default value is the FQDN of the M1 or P1 interface, depending on which interface is used for proxy traffic.

Table 14-9 Transparent Proxy Mode Authentication Settings (Continued)

Setting	Description
Surrogate Type	This setting specifies the way that transactions are associated with a user (either by IP address or using a cookie) after the user has authenticated successfully. Choose one of the following options: • **IP Address.** The appliance authenticates the user at a particular IP address. You can achieve single sign-on behavior when you choose IP-based authentication. • **Persistent Cookie.** The appliance authenticates a user on a particular application by generating a persistent cookie for each user per application. The cookie is not removed when the application is closed. • **Session Cookie.** The appliance authenticates a user on a particular application by generating a session cookie for each user per domain per application. (However, when a user provides different credentials for the same domain from the same application, the cookie is overwritten.) The cookie is removed when the application is closed. You might want to use IP-based authentication when there is only one user on a client machine and you want users to be able to achieve single sign-on behavior. You might want to choose cookie-based authentication when there are multiple users on one machine, such as a Citrix server.
Surrogate Timeout	This setting specifies how long the Web Proxy waits before asking the client for authentication credentials again. Until the Web Proxy asks for credentials again, it uses the value stored in the surrogate (IP address or cookie). It is common for user agents, such as browsers, to cache the authentication credentials so the user will not be prompted to enter credentials each time.
Secure Authentication Certificate	When Credential Encryption is enabled, you can choose whether the appliance uses the digital certificate and key shipped with the appliance or a digital certificate and key you upload here.

Figure 14-9 on page 277 shows where you configure the global authentication settings when the Web Proxy is deployed in explicit forward mode.

Figure 14-9 Explicit Forward Proxy Mode Authentication Settings

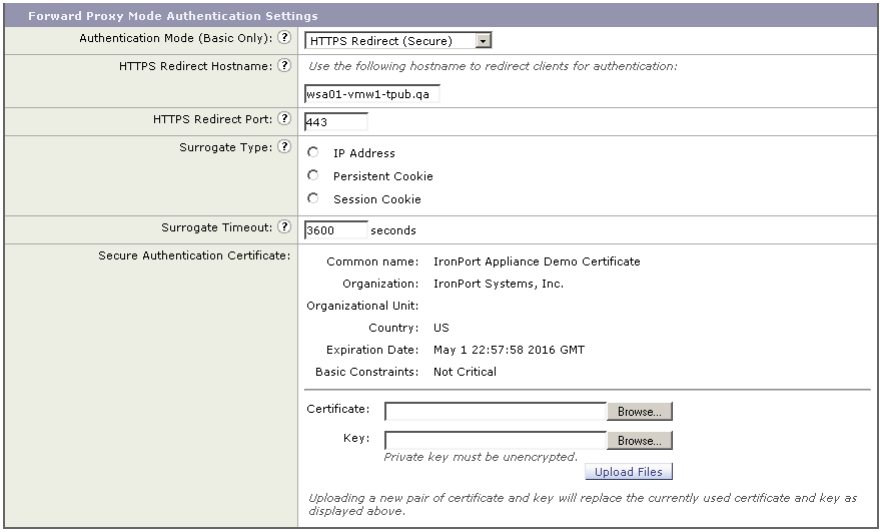

4. If the Web Proxy is deployed in explicit forward mode, edit the settings in Table 14-10.

Table 14-10 Explicit Forward Proxy Mode Authentication Settings

Setting	Description
Authentication Mode (Basic Only)	Choose whether or not the Web Proxy redirects clients to securely pass authentication credentials to the Web Proxy using HTTPS. When you choose to securely pass authentication credentials using HTTPS, additional fields appear to configure how to redirect clients. This setting applies to requests using the Basic authentication scheme only since NTLMSSP requests are secure by definition. For more information, see “Sending Authentication Credentials Securely” on page 279.

Table 14-10 Explicit Forward Proxy Mode Authentication Settings (Continued)

Setting	Description
HTTPS Redirect Hostname	Enter the short host name of the network interface on which the Web Proxy listens for incoming connections. When you enable Authentication Mode above, the Web Proxy uses this host name in the redirection URL sent to clients for authenticating users. You can enter either the following values: • **Single word host name.** You can enter the single word host name that is DNS resolvable by the client and the Web Security appliance. This allows clients to achieve true single sign-on with Internet Explorer without additional browser side setup. Be sure to enter the single word host name that is DNS resolvable by the client and the Web Security appliance. For example, if your clients are in domain `mycompany.com` and the interface on which the Web Proxy is listening has a full host name of `proxy.mycompany.com`, then you should enter `proxy` in this field. Clients perform a lookup on `proxy` and they should be able to resolve `proxy.mycompany.com`. • **Fully qualified domain name (FQDN).** You can also enter the FQDN or IP address in this field. However, if you do that and want true single sign-on for Internet Explorer and Firefox browsers, you must ensure that the FQDN or IP address is added to the client's Trusted Sites list in the client browsers. The default value is the FQDN of the M1 or P1 interface, depending on which interface is used for proxy traffic.
HTTPS Redirect Port	Enter which port to use for redirecting requests using HTTPS.

Table 14-10 Explicit Forward Proxy Mode Authentication Settings (Continued)

Setting	Description
Surrogate Type	This setting specifies the way that transactions used for authenticating the client are associated with a user (either by IP address or using a cookie) after the user has authenticated successfully. Choose one of the following options: • **IP Address.** The appliance authenticates the user at a particular IP address. You can achieve single sign-on behavior when you choose IP-based authentication. • **Persistent Cookie.** The appliance authenticates a user on a particular application by generating a persistent cookie for each user per application. The cookie is not removed when the application is closed. • **Session Cookie.** The appliance authenticates a user on a particular application by generating a session cookie for each user per application. The cookie is removed when the application is closed. You might want to use IP-based authentication when there is only one user on a client machine and you want users to be able to achieve single sign-on behavior. You might want to choose cookie-based authentication when there are multiple users on one machine, such as a Citrix server.
Surrogate Timeout	This setting specifies how long the Web Proxy waits before asking the client for authentication credentials again. Until the Web Proxy asks for credentials again, it uses the value stored in the surrogate (IP address or cookie). Note that it is common for user agents, such as browsers, to cache the authentication credentials so the user will not be prompted to enter credentials each time.
Secure Authentication Certificate	When Authentication Mode is enabled, you can choose whether the appliance uses the digital certificate and key shipped with the appliance or a digital certificate and key you upload here.

5. Submit and commit your changes.

Sending Authentication Credentials Securely

When authentication is used to identify clients using the Web, the client applications send the authentication credentials to the Web Proxy, which in turn passes them to the authentication server. How the credentials are passed from the clients to the Web Proxy depends on the authentication scheme used:

- **NTLMSSP.** The credentials are always passed to the Web Proxy securely. They are encrypted using a key specified by the Active Directory server and sent over HTTP.
- **Basic.** By default, the credentials are passed to the Web Proxy insecurely. They are encoded, but not encrypted, and sent over HTTP. However, you can configure the Web Security appliance so clients send authentication credentials securely. This works for both LDAP and NTLM Basic authentication.

When you configure the appliance to use secure client authentication for Basic authentication, the Web Proxy redirects the client back to the Web Proxy, but this time using HTTPS. The client application makes either a GET or a CONNECT request depending on how the requests are forwarded to the appliance (explicitly or transparently) and how the client application is configured to forward HTTPS requests, either using the Web Proxy or not.

Then, using the secure HTTPS connection, the clients send the authentication credentials. The appliance uses its own certificate and private key to create an HTTPS connection with the client by default. Most browsers will warn users that the certificate is not valid. To prevent users from seeing the invalid certificate message, you can upload a certificate and key pair your organization uses. When you upload a certificate and key, the private key must be *unencrypted*. For information about uploading a certificate and key, see "Installing a Server Digital Certificate" on page 411.

To configure the appliance to use secure client authentication, enable the Credential Encryption setting in the global authentication settings. For more information, see "Configuring Global Authentication Settings" on page 271. You can also use the `advancedproxyconfig > authentication` CLI command. For more information, see "Advanced Proxy Configuration" on page 71.

Accessing HTTPS Sites with Secure Client Authentication Enabled

Secure client authentication works because the Web Proxy redirects clients back to the Web Proxy using an HTTPS request. Due to this behavior, it is possible for HTTPS requests going through a Web Security appliance configured for secure client authentication to get into a redirect loop. To prevent this when enabling secure client authentication, you must also configure the appliance in one of the following ways:

- Use IP based authentication.
- Use cookie based authentication and enable HTTPS scanning on the Security Services > HTTPS Proxy page.

Note — Secure client authentication does not work with HTTPS requests in explicit forward mode.

AUTHENTICATING USING LDAP

The Lightweight Directory Access Protocol (LDAP) server database is a repository for employee directories. These directories include the names of employees along with various types of personal data such as a phone number, email address, and other information that is exclusive to the individual employee. The LDAP database is composed of objects containing attributes and values. Each object name is referred to as a distinguished name (DN). The location on the LDAP server where a search begins is called the Base Distinguished Name or base DN.

The appliance supports standard LDAP server authentication and Secure LDAP authentication. Support for LDAP allows established installations to continue using their LDAP server database to authenticate users.

For Secure LDAP, the appliance supports LDAP connections over SSL. The SSL protocol is an industry standard for ensuring confidentiality. SSL uses key encryption algorithms along with Certificate Authority (CA) signed certificates to provide the LDAP servers a way to verify the identity of the appliance.

Note — AsyncOS for Web only supports 7-bit ASCII characters for passwords when using the Basic authentication scheme. Basic authentication fails when the password contains characters that are not 7-bit ASCII.

Changing Active Directory Passwords

After Active Directory LDAP users change their account passwords, the Active Directory LDAP server authenticates them with their current or previous password, depending on the Active Directory server configuration.

If you want users to only be able to authenticate with their new password, you can reboot the Active Directory server or, you can wait for the Active Directory server to time out the old passwords.

LDAP Authentication Settings

Table 14-11 describes the authentication settings you define when you choose LDAP authentication.

Table 14-11 LDAP Authentication Settings

Setting	Description
LDAP Version	Choose the version of LDAP, and choose whether or not to use Secure LDAP. The appliance supports LDAP version 2, and LDAP version 3 software. Secure LDAP requires LDAP version 3.

Table 14-11 LDAP Authentication Settings (Continued)

Setting	Description
LDAP Server	Enter the LDAP server IP address or host name and its port number. You can specify up to three servers. The host name must be a fully-qualified domain name. For example, `ldap.example.com`. An IP address is required only if the DNS servers configured on the appliance cannot resolve the LDAP server host name. The default port number for Standard LDAP is 389. The default number for Secure LDAP is 636. If the LDAP server is an Active Directory server, enter the host name or IP address and the port of the domain controller here. Whenever possible, enter the name of the Global Catalog Server and use port 3268. However, you might want to use a local domain controller when the global catalog server is physically far away and you know you only need to authenticate users on the local domain controller. **Note:** When you configure multiple authentication servers in the realm, the appliance attempts to authorize with up to three authentication servers before failing to authenticate the transaction within that realm.
LDAP Persistent Connections (under the Advanced section)	Choose one of the following values: • **Use persistent connections (unlimited)**. Use existing connections. If no connections are available a new connection is opened. • **Use persistent connections.** Use existing connections to service the number of requests specified. When the maximum is reached, establish a new connection to the LDAP server. • **Do not use persistent connections.** Always create a new connection to the LDAP server.

Table 14-11 LDAP Authentication Settings (Continued)

Setting	Description
User Authentication	Enter values for the following fields: **Base Distinguished Name (Base DN)** The LDAP database is a tree-type directory structure and the appliance uses the Base DN to navigate to the correct location in the LDAP directory tree to begin a search. A valid Base DN filter string is composed of one or more components of the form `object-value.` For example `dc=companyname, dc=com`. **User Name Attribute** Choose one of the following values: • **uid**, **cn**, and **sAMAccountName.** Unique identifiers in the LDAP directory that specify a username. • **custom.** A custom identifier such as `UserAccount`. **User Filter Query** The User Filter Query is an LDAP search filter that locates the users Base DN. This is required if the user directory is in a hierarchy below the Base DN, or if the login name is not included in the user-specific component of that users Base DN. Choose one of the following values: • **none.** Filters any user. • **custom.** Filters a particular group of users.

Table 14-11 LDAP Authentication Settings (Continued)

Setting	Description
Query Credentials	Choose whether or not the authentication server accepts anonymous queries. If the authentication server does accept anonymous queries, choose Server Accepts Anonymous Queries. If the authentication server does not accept anonymous queries, choose Use Bind DN and then enter the following information: • **Bind DN.** The user on the external LDAP server permitted to search the LDAP directory. Typically, the bind DN should be permitted to search the entire directory. • **Password.** The password associated with the user you enter in the Bind DN field. The following text lists some example users for the Bind DN field: cn=administrator,cn=Users,dc=domain,dc=com sAMAccountName=jdoe,cn=Users,dc=domain,dc=com. If the Active Directory server is used as an LDAP server, you may also enter the Bind DN username as "DOMAIN\username."

Table 14-11 LDAP Authentication Settings (Continued)

Setting	Description
Group Authorization	Choose whether or not to enable LDAP group authorization. When you enable group authorization, enter values for the following fields: **Group Name Attribute** Choose one of the following values: • **cn.** A unique identifier in the LDAP directory that specifies the name of a group. • **custom.** A custom identifier such as `FinanceGroup.` **Group Filter Query** The Group Filter Query is an LDAP search filter that searches the LDAP directory for any matching group memberships. Choose one of the following values: • **objectclass=groupofnames** • **objectclass=groupofuniquenames** • **custom.** A custom filter such as `objectclass=person`. **Note:** The query defines the set of authentication groups which can be used in Web Security Manager policies. **Group Membership Attribute** Choose one of the following values: • **member** and **uniquemember.** Unique identifiers in the LDAP directory that specify group members. • **custom.** A custom identifier such as `UserInGroup.`

AUTHENTICATING USING NTLM

The NT Lan Manager (NTLM) authenticates users with an encrypted challenge-response sequence that occurs between the appliance and a Microsoft Windows domain controller. The NTLM challenge-response handshake occurs when a web browser attempts to connect to the appliance and before data is delivered.

When you configure an NTLM authentication realm, you do not specify the authentication scheme. Instead, you choose the scheme at the access policy group level when you configure the policy member definition. This allows you to choose different schemes for different policy groups. When you create or edit the policy group, you can choose one of the following schemes:

- Use NTLMSSP
- Use Basic or NTLMSSP
- Use Basic

Note — AsyncOS for Web only supports 7-bit ASCII characters for passwords when using the Basic authentication scheme. Basic authentication fails when the password contains characters that are not 7-bit ASCII.

Working with Multiple Active Directory Domains

AsyncOS allows you to create only one NTLM authentication realm. If your organization has multiple Active Directory domains, you can authenticate users in all domains if the following conditions exist:

- All Active Directory domains must exist in a single forest.
- There must be a trust relationship among all domains in the forest.

When you define policy group membership by group name, the web interface only displays Active Directory groups in the domain where AsyncOS created a computer account when joining the domain. To create a policy group for users in a different domain in the forest, manually enter the domain and group name in the web interface.

NTLM Authentication Settings

Table 14-12 describes the authentication settings you define when you choose NTLM authentication.

Table 14-12 NTLM Authentication Settings

Setting	Description
Active Directory Server	Enter the Active Directory server IP address or host name. You can specify up to three servers. The host name must be a fully-qualified domain name. For example, `ntlm.example.com`. An IP address is required only if the DNS servers configured on the appliance cannot resolve the Active Directory server host name. **Note:** When multiple authentication servers are configured in the realm, the appliance attempts to authorize with up to three authentication servers before failing to authorize the transaction within this realm.
Active Directory Account	Enter the following Active Directory account information: • Active Directory server domain name. • NetBIOS domain name. You only need to enter the NetBIOS domain name if the network uses NetBIOS. This field only appears when the NTLM security mode is set to "domain" using the `setntlmsecuritymode` CLI command. • Computer account location. **Note:** You must click **Join Domain** to enter an Active Directory username and password. For more information about entering the Active Directory account information, see "Joining the Active Directory Domain" on page 288.
Join Domain button (Active Directory User)	When you click **Join Domain**, enter the name and password for the Active Directory user. If the appliance and the Active Directory server are in the same domain, any valid user that is a member of User Domain is allowed. However, depending on the Active Directory server configuration, this user might need Domain Admin Group or Enterprise Admin Group credentials. For example: • If the appliance and the Active Directory server are not in the same domain, the Active Directory user must be a member of the Domain Admin Group. • If the Active Directory server configuration is a forest, the Active Directory user must be a member of the Enterprise Admin Group.

Table 14-12 NTLM Authentication Settings (Continued)

Setting	Description
Network Security	Configure whether or not the Active Directory server is configured to require signing. When you enable this check box, the appliance uses Transport Layer Security (TLS) when communicating with the Active Directory server.

Joining the Active Directory Domain

When you configure an NTLM realm, you must enter information to join the Active Directory domain to set up a computer account in the domain. An Active Directory computer account is an account that uniquely identifies the computer on the domain. It is also referred to as a machine trust account.

After you enter the Active Directory account information in the authentication realm, click the **Join Domain** button to set up a computer account. Use the Location field to define the organizational directory where AsyncOS should create the computer account in the Active Directory domain.

Figure 14-10 on page 289 shows where you join an Active Directory domain.

Figure 14-10 Joining an Active Directory Domain

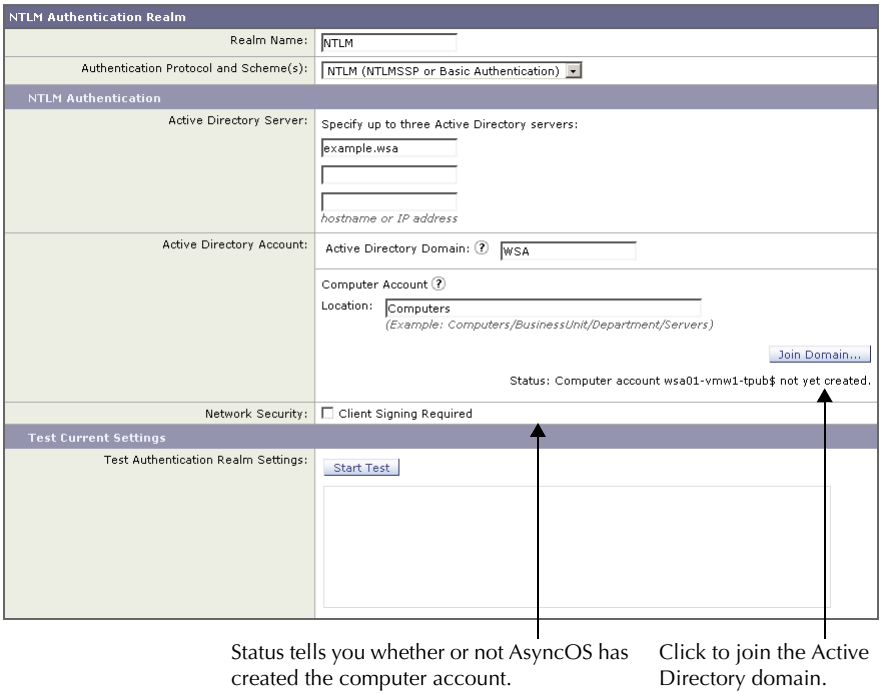

When you click **Join Domain**, you are prompted to enter login credentials for the Active Directory server. The login information is used only to create the Active Directory computer account and is not saved. Enter the login information and click **Create Account**.

Note — You must enter the sAMAccountName user name for the Active Directory user. Also, verify that users enter their sAMAccountName user name when they log in to their computers.

Once an account is created, the status of the account creation is displayed below the Join Domain button. If the account creation fails, the status and reason for error is displayed.

Also, when you view all realms on the Network > Authentication page, the appliance displays warning text in red saying that the domain was not joined for any realm that did not create a computer account.

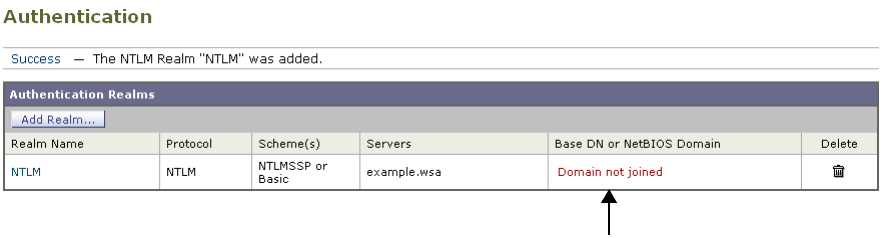

Red text indicates that the domain was not joined and no computer account was created.

AsyncOS only creates an Active Directory computer account when you edit the authentication realm Active Directory information or when the appliance reboots.

Note — To successfully join the Active Directory domain, the time difference between the Web Security appliance and the Active Directory server should be less than the time specified in the "Maximum tolerance for computer clock synchronization" option on the Active Directory server. When you use Network Time Protocol (NTP) to specify the current time on the Web Security appliance, remember that the default time server is time.ironport.com. This may affect the time difference between the appliance and the Active Directory server.

SUPPORTED AUTHENTICATION CHARACTERS

This section lists the characters the Web Security appliance supports when it communicates with LDAP and Active Directory servers. For authentication to work properly, verify that your authentication servers only use the supported characters listed in this section.

For example, according to Table 14-13, the appliance can validate users with the following Active Directory user name:

```
jsmith#123
```

And according to Table 14-13, the appliance cannot validate users with the following Active Directory user name:

```
jsmith+
```

Active Directory Server Supported Characters

Table 14-13 lists the characters the Web Security appliance supports for the User Name field for Active Directory servers.

Table 14-13 Supported Active Directory Server Characters — User Name Field

Supported Characters	Characters Not Supported
A...Z a...z 0 1 2 3 4 5 6 7 8 9 ` ~ ! # $ % ^ & ( ) _ - { } ' . space	/ \ [ ] : ; | = , + * ? < > @ "

Note — The Web Security appliance supports the percent (%) character for end users browsing the web. However, you cannot use a user name with the percent (%) character to join the Active Directory domain when you create an NTLM authentication realm.

Table 14-14 lists the characters the Web Security appliance supports for the Password field for Active Directory servers.

Table 14-14 Supported Active Directory Server Characters — Password Field

Supported Characters	Characters Not Supported
A...Z a...z 0 1 2 3 4 5 6 7 8 9 ` ~ ! # $ ^ & ( ) _ - { } ' . / [ ] : | * ? @ + \ , ; " = < > space	N/A

Table 14-15 lists the characters the Web Security appliance supports for the Location field for Active Directory servers. You enter the location string in the Location field when you configure an NTLM authentication realm.

Table 14-15 Supported Active Directory Server Characters — Location Field

Supported Characters	Characters Not Supported
A...Z a...z 0 1 2 3 4 5 6 7 8 9 ` ~ ! # $ ^ & ( ) _ - { } ' . / [ ] : | * ? @ space	+ \ , ; " = < > **Note:** The appliance does not support these characters even when they are escaped with a backslash (\ ) character.

Table 14-16 lists the characters the Web Security appliance supports for the Group field for Active Directory servers.

Table 14-16 Supported Active Directory Server Characters — Group Field

Supported Characters	Characters Not Supported
A...Z a...z 0 1 2 3 4 5 6 7 8 9 ` ~ ! # $ % ^ & ( ) _ - { } ' . @ space	/ \ [ ] : ; | = , + * ? < > "

Note — You can only use the backslash (\ ) character as a separator between the domain name and a user or group name, or as a separator between organizational units (OU) in the location string for an Active Directory server. You cannot use it as part of a domain name, user name, group name, or location name.

LDAP Server Supported Characters

Table 14-17 lists the characters the Web Security appliance supports for the User Name field for LDAP servers.

Table 14-17 Supported LDAP Server Characters — User Name Field

Supported Characters	Characters Not Supported
A...Z a...z 0 1 2 3 4 5 6 7 8 9 ` ~ ! # $ % ^ & ( ) _ - { } ' . **Note:** The appliance only supports the '(' and ')' characters when they are escaped with a backslash (\ ) character.	/ \ [ ] : ; | = , + * ? < > @ "

Table 14-18 lists the characters the Web Security appliance supports for the Password field for LDAP servers.

Table 14-18 Supported LDAP Server Characters — Password Field

Supported Characters	Characters Not Supported
A...Z a...z 0 1 2 3 4 5 6 7 8 9 ` ~ ! # $ % ^ & ( ) _ - { } @ ' . / \ [ ] : | = * ? < > " , ; + space	N/A

Table 14-19 lists the characters the Web Security appliance supports for the Group field for LDAP servers.

Table 14-19 Supported LDAP Server Characters — Group Field

Supported Characters	Characters Not Supported
A...Z a...z 0 1 2 3 4 5 6 7 8 9 ` ~ ! # $ % ^ & ( ) _ - { } @ ' . / \ [ ] : | = * ? < > " space **Note:** The appliance only supports the '(' and ')' characters when they are escaped with a backslash (\ ) character.	, ; +

Table 14-20 lists the characters the Web Security appliance supports for the Custom User Filter Query Field field for LDAP servers.

Table 14-20 Supported LDAP Server Characters — Custom User Filter Query Field

Supported Characters	Characters Not Supported
A...Z a...z 0 1 2 3 4 5 6 7 8 9 ` ~ ! # $ % ^ & ( ) _ - { } ' . space	@ / \ [ ] : | = * ? < > " , ; +

Table 14-21 lists the characters the Web Security appliance supports for the Custom Group Filter Query Field field for LDAP servers.

Table 14-21 Supported LDAP Server Characters — Custom Group Filter Query Field

Supported Characters	Characters Not Supported
A...Z a...z 0 1 2 3 4 5 6 7 8 9 ` ~ ! # $ % ^ & ( ) _ - { } @ ' . / \ [ ] : | = * ? < > " space	, ; +

CHAPTER 15

L4 Traffic Monitor

This chapter contains the following information:

ABOUT L4 TRAFFIC MONITOR

The Web Security appliance has an integrated Layer-4 Traffic Monitor that detects rogue traffic across all network ports and stops malware attempts to bypass port 80. Additionally, when internal clients are infected with malware and attempt to phone-home across non-standard ports and protocols, the L4 Traffic Monitor prevents phone-home activity from going outside the corporate network.

HOW THE L4 TRAFFIC MONITOR WORKS

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

All web destinations fall under one of the following categories:

- **Known allowed address.** Any IP address or host name listed in the Allow List property. These addresses appear in the log files as "whitelist" addresses.
- **Unlisted address.** Any IP address that is not known to be a malware site nor is a known allowed address. They are not listed on the Allow List or Additional Suspected Malware Addresses properties, nor are they listed in the L4 Traffic Monitor Database as a known malware site. These addresses do not appear in the log files.
- **Ambiguous address.** These addresses appear in the log files as "greylist" addresses. They include any of the following addresses:
  - Any *IP address* that is associated with both an unlisted *host name* and a known malware *host name*.
  - Any *IP address* that is associated with both an unlisted *host name* and a *host name* from the Additional Suspected Malware Addresses property.
- **Known malware address.** These addresses appear in the log files as "blacklist" addresses. They include any of the following addresses:
  - Any IP address or host name that the L4 Traffic Monitor Database determines to be a known malware site and *not* listed in the Allow List.
  - Any *IP address* that is listed in the Additional Suspected Malware Addresses property and *not* listed in the Allow List and *not* determined to be ambiguous.

Note — You can define the Allow List and the Additional Suspected Malware Addresses properties on the Web Security Manager > L4 Traffic Monitor Policies page.

The L4 Traffic Monitor listens to and monitors network ports for rogue activity. It performs one of the following actions on all traffic on your network:

- **Allow.** It always allows traffic to and from known allowed and unlisted addresses.
- **Monitor.** It monitors traffic under the following circumstances:
  - When the Action for Suspected Malware Addresses option is set to Monitor, it always monitors all traffic that is not to or from a known allowed address.
  - When the Action for Suspected Malware Addresses option is set to Block, it monitors traffic to and from ambiguous addresses.
- **Block.** When the Action for Suspected Malware Addresses option is set to Block, it blocks traffic to and from known malware addresses.

The L4 Traffic Monitor Database

The L4 Traffic Monitor uses and maintains its own internal database. This database is continuously updated with matched results for IP addresses and domain names. Additionally, the database table receives periodic updates from the IronPort update server at the following location:

`https://update-manifests.ironport.com`

For information about update intervals and the IronPort update server, see "Component Updates" on page 391.

CONFIGURING THE L4 TRAFFIC MONITOR

The L4 Traffic Monitor can be enabled as part of an initial system setup using the System Setup Wizard. By default, the L4 Traffic Monitor is enabled and set to monitor traffic on all ports. This includes DNS and other services.

Note — To monitor true client IP addresses, the L4 Traffic Monitor should always be configured inside the firewall and before network address translation (NAT). For more information about deploying the L4 Traffic Monitor, see "Deploying the L4 Traffic Monitor" on page 27.

You can configure the following settings:

- **Global L4 Traffic Monitor settings.** You can enable or disable the L4 Traffic Monitor after an initial configuration and configure which TCP ports to monitor. Use the Security Services > L4 Traffic Monitor page. For more information see "Configuring L4 Traffic Monitor Global Settings" on page 299.
- **L4 Traffic Monitor policies.** When the L4 Traffic Monitor is enabled, you configure specific policies for managing traffic. Use the Web Security Manager > L4 Traffic Monitor Policies page. For more information see "Configuring L4 Traffic Monitor Policies" on page 300.

Configuring L4 Traffic Monitor Global Settings

On the Security Services > L4 Traffic Monitor page, you can configure the L4 Traffic Monitor global settings and update the L4 Traffic Monitor anti-malware rules.

Figure 15-1 Security Services > L4 Traffic Monitor Page

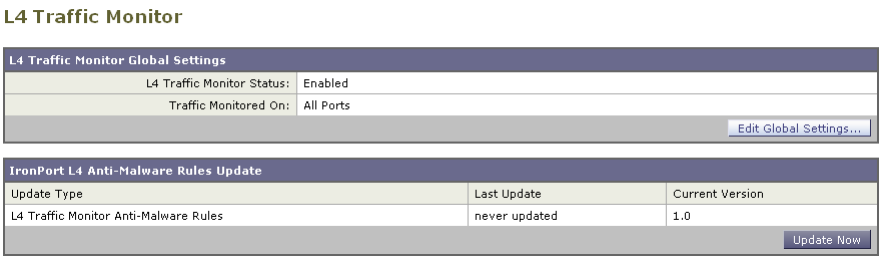

To configure L4 Traffic Monitor global settings:

1. Navigate to the Security Services > L4 Traffic Monitor page.
2. Click Edit Global Settings.
3. Choose whether or not to enable the L4 Traffic Monitor.
4. When you enable the L4 Traffic Monitor, choose which ports it should monitor:
    - **All ports.** Monitors all 65535 TCP ports for rogue activity.

- **All ports except proxy ports.** Monitors all TCP ports except the following ports for for rogue activity.
  - Ports configured in the "HTTP Ports to Proxy" property on the Security Services > Proxy Settings page (usually port 80).
  - Ports configured in the "Transparent HTTPS Ports to Proxy" property on the Security Services > HTTPS Proxy page (usually port 443).

5. Submit and commit the changes.

Updating L4 Traffic Monitor Anti-Malware Rules

To update the L4 Traffic Monitor anti-malware rules:

1. Navigate to the Security Services > L4 Traffic Monitor page.
2. Click **Update Now**.

   The Web Security appliance contacts the component update server and updates the L4 Traffic Monitor anti-malware rules. For more information about the component update server, see "Component Updates" on page 391.

Configuring L4 Traffic Monitor Policies

When the L4 Traffic Monitor is enabled, you can configure how it should manage traffic over the configured TCP ports. It can perform the following actions on traffic over the TCP ports:

- Allow
- Monitor
- Block

For more information about how the L4 Traffic Monitor handles traffic, see "How the L4 Traffic Monitor Works" on page 297.

The actions the L4 Traffic Monitor takes depends on the L4 Traffic Monitor policies you configure.

To configure L4 Traffic Monitor policies:

1. Navigate to the Web Security Manager > L4 Traffic Monitor page.
2. Click **Edit Settings**.

3. On the Edit L4 Traffic Monitor Policies page, configure the L4 Traffic Monitor policies described in Table 15-1.

Table 15-1 L4 Traffic Monitor Policies

Property	Description
Allow List	Enter zero or more address to which the L4 Traffic Monitor should always allow clients to connect. Separate multiple entries with a space or comma. For a list of valid address formats you can use, see "Valid Formats" on page 302. **Note:** Entering a domain name such as example.com also matches www.example.com and hostname.example.com. Connections to all destinations in this list are always allowed and the traffic is not logged. The appliance does not check the destinations against the L4 Traffic Monitor anti-malware rules or the additional suspected malware addresses listed on the same page. For example, if IP address 10.1.1.1 appears in both the Allow List and the Additional Suspected Malware Addresses fields, then the L4 Traffic Monitor always allows requests for 10.1.1.1.
Actions for Suspected Malware Addresses	Choose whether to monitor or block traffic destined for a known malware address. For a definition of known malware address, see "How the L4 Traffic Monitor Works" on page 297. • **Monitor.** Scans all traffic for domains and IP addresses that match entries in the L4 Traffic Monitor database. The Monitor option does not block suspicious traffic. This setting is useful for identifying infected clients without affecting the user experience. • **Block.** Scans all traffic for domains and IP addresses that match entries in the appliance administrative lists and the block list database and then blocks any traffic it finds. This setting is useful for identifying infected clients and stopping malware attempts through non-standard ports. When you choose to block suspected malware traffic, you can also choose whether or not to always block ambiguous addresses. By default, ambiguous addresses are monitored. For a definition of ambiguous address, see "How the L4 Traffic Monitor Works" on page 297.

Table 15-1 L4 Traffic Monitor Policies

Property	Description
Additional Suspected Malware Addresses (optional)	Enter zero or more known addresses that the L4 Traffic Monitor should consider as a possible malware. For a list of valid address formats you can use, see "Valid Formats" on page 302. If you choose to block suspected malware addresses, the L4 Traffic Monitor will either block or monitor these addresses depending on whether it determines them to be known malware addresses or ambiguous addresses. For definitions of ambiguous and known malware addresses, see "How the L4 Traffic Monitor Works" on page 297. If you choose to monitor suspected malware addresses, it will monitor these addresses. **Note:** Adding internal IP addresses to the Additional Suspected Malware Addresses list causes legitimate destination URLs to show up as malware in L4 Traffic Monitor reports. To avoid this type of erroneous reporting, do not enter internal IP addresses in the "Additional Suspected Malware Addresses" field on the Web Security Manager > L4 Traffic Monitor Policies page.

Note — If the L4 Traffic Monitor is configured to block, the L4 Traffic Monitor and the Web Proxy must be configured on the same network. Use the Network > Routes page to confirm that all clients are accessible on routes that are configured for data traffic.

4. Submit and commit your changes.

Valid Formats

When you add addresses to the Allow List or Additional Suspected Malware Addresses properties, separate multiple entries with whitespace or commas. You can enter addresses in any of the following formats:

- **IP address.** For example, 10.1.1.0.
- **CIDR address.** For example, 10.1.1.0/24.
- **Domain name.** For example, example.com. Entering a domain name such as example.com will also match www.example.com and hostname.example.com.
- **Hostname.** For example, crm.example.com.

VIEWING L4 TRAFFIC MONITOR ACTIVITY

The S-Series appliance supports several options for generating feature specific reports and interactive displays of summary statistics.

Reports

You can use options on the Monitor > Reports pages of the web interface to select a type of report, capture data, schedule periodic email delivery, and archive reports.

For more information about working with reports, see "Reporting Overview" on page 324.

Monitoring Activity and Viewing Summary Statistics

The Monitor > L4 Traffic Monitor page provides statistical summaries of monitoring activity. You can interactively update these displays by specifying a time range of hour, day, week or month. Additionally, you have the option to print these display pages and export the raw data to a file.

You can use the following displays and reporting tools to view the results of L4 Traffic Monitor activity:

Table 15-2 L4 Traffic Monitor Scanning Data

To view...	See...
Client statistics	Monitor > Client Activity
Malware statistics Port statistics	Monitor > L4 Traffic Monitor
L4 Traffic Monitor log files	System Administration > Log Subscriptions • trafmon_errlogs • trafmonlogs

Note — If the Web Proxy is configured as a forward proxy and L4 Traffic Monitor is set to monitor all ports, the IP address of the proxy's data port is recorded and displayed as a client IP address in the client activity report on the Monitor > Client Activity page. If the Web Proxy is configured as a transparent proxy, enable IP spoofing to correctly record and display the client IP addresses.

L4 Traffic Monitor Log File Entries

The L4 Traffic Monitor log file provides a detailed record of monitoring activity. For more information about the L4 Traffic Monitor log, see "Traffic Monitor Log" on page 355.

CHAPTER 16

Monitoring

This chapter contains the following information:

MONITORING SYSTEM ACTIVITY

Administrators and executive management require information to better understand evolving corporate threats. While the Web Security appliance controls the malware threat to a corporate environment, comprehensive monitoring and reporting tools provide insight to threats that are monitored or blocked, and display actionable data such as top clients infected to help you manage the presence of malware.

The chapter introduces you to the monitoring tools you can use to monitor system activity and help you interpret data specific to each Web Security appliance security component. The Monitor tab contains a collection of system data and graphical displays for the following types of information:

- **Security Services** — Summary displays of transaction data derived from the results of filtering policies.
- **Suspect Transactions Detected** — Summary charts that represent the percentages of traffic that was blocked by S-Series filtering and scanning features.
- **Top Sites by Malware** — Categorical displays of monitored and blocked transactions to web sites containing malware.
- **High-Risk and Malware Activity** — Summary displays of client malware activity and high-risk web sites.

Note — You can also use appliance reports to monitor appliance activity. For more information about creating and using reports, see “Reporting” on page 323.

USING THE MONITOR TAB

The Monitor tab provides several options for viewing system data. This section describes those options and explains the information displayed on each of the following pages: Overview, L4 Traffic Monitor, Client Web Activity, Client Malware Risk, Web Site Activity, Anti-Malware, URL Categories, and Web Reputation Filters.

Monitor tab display pages provide a colorful overview of system activity and support multiple options for viewing system data. For example, you can update and sort data to provide real-time visibility into resource utilization and web traffic trouble spots. You can also search each page for web site and client-specific data.

Changing the Timeframe

You can update the data displayed for each security component using the Time-Range field. This option allows you to generate updates for the last hour, day, week, or 30-day period. For example:

Figure 16-1 Selecting Data Time Range

Report data is displayed as follows:

Table 16-1 Time Intervals for Data Collection

For this time increment…	Data is returned in…
Hour	Sixty (60) complete minutes plus up to 5 additional minutes.
Day	One hour intervals for the last 24 hours and including the current partial hour.
Week	One day intervals for the last 7 days plus the current partial day.
Month (30 days)	One day intervals for the last 30 days plus the current partial day.

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.

Searching Data

The Search option at the bottom of each display page returns data for a particular web site or client.

Figure 16-2 Searching for Web Sites or Clients

You can search for an exact match of a web site, client IP address or user ID, or you can search for web sites or clients that start with a specific text string.

Note — You need to configure authentication to view client user IDs instead of client IP addresses.

OVERVIEW PAGE

The Monitor > Overview page displays the Overview report. This report contains highlights of the System Status report and provides summary system traffic and security risk summary data.

The following sections appear on the Overview report:

Overview Report Section	Description
System Overview	Displays the following information: • Web Proxy traffic characteristics, including per-minute data averages for client transactions, system bandwidth, average system response time, and total current connections. • System resource utilization statistics. Use the System Status Details link to access the full System Status report.
Total Web Activity	Compares total clean transactions and total suspect transactions in a trend graph over time. Suspect transactions include all monitored and blocked transactions.
Suspect Transactions Detected	Compares various types of suspect transactions in a trend graph over time.
Security Services Summary	Displays the following information: • Web Proxy total suspect transactions and suspect transactions by type. Suspect transactions include requests blocked URL category, requests blocked by Web Reputation Filters, transactions detected by Anti-Malware scanning, and other blocked transactions. "Other blocked transactions" includes transactions blocked by various policy settings such as the file size limit. • L4 Traffic Monitor suspect connections. For information about the specific monitored and blocked transactions, see "L4 Traffic Monitor Data Page" on page 310.
Top URL Categories	Lists the top 10 URL categories matched for the specified time range. For more information, see "URL Categories Page" on page 314.
Top Malware Categories	Lists the top 10 malware categories detected for the specified time range. For more information, see "Anti-Malware Page" on page 313.

Export links that are visible on each page, are used to export raw data. For more information, see "Exporting Report Data" on page 329.

L4 TRAFFIC MONITOR DATA PAGE

The Monitor > L4 Traffic Monitor page displays information about malware ports and malware sites that the L4 Traffic Monitor 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 information about malware detected by the Anti-Malware DVS engine, see "Anti-Malware Page" on page 313.

CLIENTS PAGES

You can use the following pages to monitor client activity:

- Monitor > Clients > Web Activity page — This page shows the Client Web Activity report, which includes the following information:
  - Top clients by total web transactions
  - Top clients by blocked web transactions

  The client details table provides additional details including, bandwidth usage and the amount of bandwidth saved by blocking.

  The user ID's and client IP addresses are interactive and link to a Client Detail page that provides information respective to each client.
- Monitor > Clients > Malware Risk page — This page shows the Client Malware Risk report, which includes the following information:
  - Web Proxy top clients by malware risk (number of transactions)
  - L4 Traffic Monitor top clients by malware risk (number of connections)

  The client details at the bottom of the page display the same data as the graphs, but for *all* clients and in table format. In addition, the All tab for Web Proxy transactions provides information about the bandwidth that was saved by blocking, and it shows how many monitored and blocked malware transactions were detected at request time or detected at response time.

  The user ID's and client IP addresses are interactive and link to a Client Detail page that provides detailed information respective to each client.
- Client Detail page — This page shows all the web activity and malware risk data for a particular client during the specified time range. It includes the following information:
  - Completed and blocked web transactions
  - Web Proxy monitored and blocked malware transactions
  - L4 Traffic Monitor malware connections
  - URL categories matched
  - Malware threats detected
  - Suspect user agents detected

Note — The client reports sometimes show a user with an asterisk (*) at the end of the user name. For example, the Client Web Activity 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.

WEB SITE ACTIVITY PAGE

Use the following pages to monitor high-risk web sites accessed during a specific time range:

- Monitor > Web Site Activity page — This page shows the Web Site Activity report, which includes the following information:
  - Top five sites by high-risk transactions detected. A high-risk transaction is any monitored or blocked transaction.
  - Top five sites by malware transactions detected.

  The site details section at the bottom of the page lists all of the sites with high-risk transactions. You can use column headings to sort the data and each URL links to the Web Site Detail page.
- Web Site Detail page — This page shows the high-risk transactions for the site in a trend graph that uses a different color for each type of high-risk transaction.

  The Summary tab shows the same information as the trend graph, but in table format. It shows the transactions blocked by URL filtering, transactions blocked by Web Reputation Filters, transactions detected by Anti-Malware scanning, other blocked transactions, total high-risk transactions, and URL categories of the site. The All tab displays bandwidth saved by blocking and includes detail about transactions detected by Anti-Malware scanning.

  The Other Blocked Transactions column displays transactions blocked by a policy rule. This data includes the following conditions:

  - File size over limit
  - File type not allowed
  - User agent not allowed
  - Protocol not allowed
  - Authentication denied
  - Attempted HTTP tunneling (CONNECT) on disabled port

  User agents blocked by a policy configuration are recorded as “other blocked transactions.” Suspect user agents detected by the Anti-Malware DVS engine are recorded as blocked by Anti-Malware scanning.

ANTI-MALWARE PAGE

Use the following pages to monitor malware detected by the Anti-Malware DVS engine:

- Monitor > Anti-Malware page — This page shows the Anti-Malware report, which includes the following information:
  - Top malware categories detected (by number of transactions)
  - Top malware threats detected (by number of transactions)

  The Malware Categories and Malware Threats sections show the same data as the graphs, but in table format. In addition, these sections include information about the bandwidth saved by blocking.
- Malware Category page — The name of a malware category on the Anti-Malware report is a link to the Malware Category page. The Malware Category report shows detailed information about a particular malware category. The trend graph at the top of the report shows the monitored and blocked transactions for the category during the specified time range. The table at the bottom lists the detected malware threats that belong to the malware category and shows the number of monitored and blocked transactions for each.
- Malware Threat page — The name of a malware threat on either the Anti-Malware report or the Malware Category report is a link to the Malware Threat page. The Malware Threat 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.

URL CATEGORIES PAGE

Use the Monitor > URL Categories page to view the URL Categories report. This report shows the top 10 URL categories by completed transactions and the top 10 URL categories by blocked transactions for a specified time range. Competed transactions include both clean transactions and monitored transactions.

The URL Categories Matched section shows all matched categories during a specified time range for both completed and blocked transactions. You can use column headings to sort data, and the Items Displayed menu changes the number of URL categories displayed in the list.

The percentage of uncategorized URLs in the URL Categories report is typically around 15-20%. If the percentage of uncategorized URLs is higher than that, consider the following options:

- For specific localized URLs, you can create custom URL Categories and apply them on to specific users or group policies. For more information about custom URL Categories, see "Custom URL Categories" on page 216.
- You can report misclassified and uncategorized URLs to the IronPort support portal at the following URL:
  `https://supportportal.ironport.com/irppcnctr/srvcd?u=http://secure-support.soma.ironport.com/subproducts/s_series&sid=900019#`
  These get picked up and get evaluated for subsequent rule updates.
- Verify Reputation Filtering and Anti-Malware Filtering is enabled. Often times, the correlation between malware and URLs with suspect content is high and it is likely that they may get caught by subsequent filters. The system pipeline is set up to catch malicious traffic with other downstream filters if URL filtering does not have a verdict.

WEB REPUTATION FILTERS PAGE

Use the Monitor > Web Reputation Filters page to view the Web Reputation Filters report. This report shows the result of Web Reputation filtering for transactions during a specified time range.

The Web Reputation Actions trend graph compares the following types of web reputation actions during the specified time range:

- Block
- Scan Further: Malware Detected
- Scan Further: Clean
- Allow

The Web Reputation Actions (Volume) section displays the data as percentages in table format.

The Web Reputation Filters report also shows the configured score ranges for the Block, Scan Further, and Allow actions. In addition, it displays a breakdown by score for each filtered transaction.

Note — If the result of Web Reputation filtering is to Scan Further, the transaction is passed to the Anti-Malware DVS engine for additional scanning.

SYSTEM STATUS PAGE

Use the Monitor > System Status page to monitor the System Status. This page displays the current status and configuration of the Web Security appliance. The following table describes each display.

Table 16-2System Status

This Section...	Displays
Web Security Appliance Status	• System uptime • System resource utilization — CPU usage, RAM usage, and percentage of disk space used for reporting and logging.
Proxy Traffic Characteristics	• Transactions per second • Bandwidth • Response time • Cache hits • Connections
Current Configuration	Web Proxy settings: • Web Proxy Status — enabled or disabled. • Deployment Topology. • Web Proxy Mode — forward or transparent. • IP Spoofing — enabled or disabled. L4 Traffic Monitor settings: • L4 Traffic Monitor Status — enabled or disabled. • L4 Traffic Monitor Wiring. • L4 Traffic Monitor Action — monitor or block.

SNMP MONITORING

The IronPort AsyncOS operating system supports system status monitoring via SNMP (Simple Network Management Protocol). This includes IronPort's Enterprise MIB, asyncoswebsecurityappliance-mib.txt. The asyncoswebsecurityappliance-mib helps administrators better monitor system health. In addition, this release implements a read-only subset of MIB-II as defined in RFCs 1213 and 1907. (For more information about SNMP, see RFCs 1065, 1066, and 1067.) Please note:

- SNMP is **off** by default.
- SNMP SET operations (configuration) are not implemented.
- AsyncOS supports SNMPv1, v2, and v3.
- The use of SNMPv3 with password authentication and DES Encryption is mandatory to enable this service. (For more information on SNMPv3, see RFCs 2571-2575.) You are required to set a SNMPv3 passphrase of at least 8 characters to enable SNMP system status monitoring. The first time you enter a SNMPv3 passphrase, you must re-enter it to confirm. The `snmpconfig` command "remembers" this phrase the next time you run the command.
- The SNMPv3 username is: v3get.

```
> snmpwalk -v 3 -l AuthNoPriv -u v3get -a MD5 ironport serv.example.com
```

- If you use only SNMPv1 or SNMPv2, you must set a community string. The community string does not default to `public`.
- For SNMPv1 and SNMPv2, you must specify a network from which SNMP GET requests are accepted.
- To use traps, an SNMP manager (not included in AsyncOS) must be running and its IP address entered as the trap target. (You can use a hostname, but if you do, traps will only work if DNS is working.)

Use the `snmpconfig` command to configure SNMP system status for the appliance. After you choose and configure values for an interface, the appliance responds to SNMPv3 GET requests. These version 3 requests must include a matching password. By default, version 1 and 2 requests are rejected. If enabled, version 1 and 2 requests must have a matching community string.

MIB Files

IronPort Systems provides "enterprise" MIBs for Email and Web Security appliances as well as a "Structure of Management Information" (SMI) file:

- asyncoswebsecurityappliance-mib.txt — an SNMPv2 compatible description of the Enterprise MIB for IronPort Web Security appliances.

- ASYNCOS-MAIL-MIB.txt — an SNMPv2 compatible description of the Enterprise MIB for IronPort Email Security appliances.
- IRONPORT-SMI.txt — defines the role of the asyncoswebsecurityappliance-mib in IronPort's SNMP managed products.

These files are available on the documentation CD included with your IronPort appliance. You can also find these files on the IronPort Customer Support portal.

Hardware Objects

Hardware sensors conforming to the Intelligent Platform Management Interface Specification (IPMI) report temperature, fan speed, and power supply status.

Table 16-3 shows what hardware derived objects are available for monitoring on what models. The number displayed is the number of instances of that object that can be monitored. For example, you can query the RPMs for 4 fans in the S350 appliance.

Table 16-3 Number of Hardware Objects per IronPort Appliance

Model	Ambient Temp	Fans	Power Supply	Disk Status	NIC Link
S160	1	2	1	2	6
S350	1	4	2	6	6
S360	1	4	2	4	6
S650	1	4	2	6	6
S660	1	4	2	6	6

Hardware Traps

Table 16-4 lists the temperature and hardware conditions that cause a hardware trap to be sent:

Table 16-4 Hardware Traps: Temperature and Hardware Conditions

Model	High Temp (Ambient)	Fan Failure	Power Supply	RAID	Link
S160/ S350/ S360/ S650/ S660	47C	0 RPMs	Status Change	Status Change	Status Change

Status change traps are sent when the status changes. Fan Failure and high temperature traps are sent every 5 seconds. The other traps are failure condition alarm traps — they are sent once when the state changes (healthy to failure). It is a good idea to poll for the hardware status tables and identify possible hardware failures before they become critical. Temperatures within 10 per cent of the critical value may be a cause for concern.

Note that failure condition alarm traps represent a critical failure of the individual component, but may not cause a total system failure. For example, a single fan or power supply can fail on a S650 appliance and the appliance will continue to operate.

SNMP Traps

SNMP provides the ability to send traps, or notifications, to advise an administration application (an SNMP management console, typically) when one or more conditions have been met. Traps are network packets that contain data relating to a component of the system sending the trap. Traps are generated when a condition has been met on the SNMP agent (in this case, the IronPort appliance). After the condition has been met, the SNMP agent then forms an SNMP packet and sends it over port 162, the standard SNMP trap port. In the example below, the trap target of `10.1.1.29` and the Trap Community string are entered. This is the host running the SNMP management console software that will receive the SNMP traps from the IronPort appliance.

You can configure SNMP traps (enable or disable specific traps) when you enable SNMP for an interface. To specify multiple trap targets: when prompted for the trap target, you may enter up to 10 comma separated IP addresses.

CLI Example

In the following example, the `snmpconfig` command is used to enable SNMP on the "PublicNet" interface on port 161. A passphrase for version 3 is entered and then re-entered for confirmation. The system is configured to service version 1 and 2 requests, and the community string `public` is entered for GET requests from those versions 1 and 2. The trap target of `10.1.1.29` is entered. Finally, system location and contact information is entered.

```
example.com> snmpconfig
Current SNMP settings:
SNMP Disabled.

Choose the operation you want to perform:
- SETUP - Configure SNMP.
[]> setup

Do you want to enable SNMP? [N]> y

Please choose an IP interface for SNMP requests.
1. Management (192.168.1.1/24: wsa01-vmw1-tpub.qa)
[1]>

Enter the SNMPv3 passphrase.
```

```
>
Please enter the SNMPv3 passphrase again to confirm.
>
Which port shall the SNMP daemon listen on?
[161]>

Service SNMP V1/V2c requests? [N]> y

Enter the SNMP V1/V2c community string.
[]> public

From which network shall SNMP V1/V2c requests be allowed?
[192.168.1.1]>

Enter the Trap target as a host name, IP address or list of IP
addresses separated by commas (IP address preferred). Enter “None” to
disable traps.
[None]> 10.1.1.29

Enter the Trap Community string.
[]> tcomm

Enterprise Trap Status
1. CPUUtilizationExceeded       Disabled
2. RAIDStatusChange             Enabled
3. connectivityFailure          Disabled
4. fanFailure                   Enabled
5. highTemperature              Enabled
6. keyExpiration                Enabled
7. linkDown                     Enabled
8. linkUp                       Enabled
9. memoryUtilizationExceeded    Disabled
10. powerSupplyStatusChange     Enabled
11. resourceConservationMode    Enabled
12. updateFailure               Enabled
13. upstream_proxy_failure      Enabled

Do you want to change any of these settings? [N]> y

Do you want to disable any of these traps? [Y]> n

Do you want to enable any of these traps? [Y]> y

Enter number or numbers of traps to enable. Separate multiple numbers
with commas.
[]> 1,3
```

```
What threshold would you like to set for CPU utilization?
[95]>

What URL would you like to check for connectivity failure?
[http://downloads.ironport.com]>

Enterprise Trap Status
1. CPUUtilizationExceeded       Enabled
2. RAIDStatusChange             Enabled
3. connectivityFailure          Enabled
4. fanFailure                   Enabled
5. highTemperature              Enabled
6. keyExpiration                Enabled
7. linkDown                     Enabled
8. linkUp                       Enabled
9. memoryUtilizationExceeded    Disabled
10. powerSupplyStatusChange     Enabled
11. resourceConservationMode    Enabled
12. updateFailure               Enabled
13. upstream_proxy_failure      Enabled
Do you want to change any of these settings? [N]>

Enter the System Location string.
[Unknown: Not Yet Configured]> Network Operations Center - west; rack
#30, position 3

Enter the System Contact string.
[snmp@localhost]> Joe Administrator, x8888

Current SNMP settings:
Listening on interface “Management” 192.168.1.1 port 161.
SNMP v3: Enabled.
SNMP v1/v2: Enabled, accepting requests from subnet 192.168.1.1.
SNMP v1/v2 Community String: public
Trap target: 10.1.1.29
Location: Network Operations Center - west; rack #30, position 3
System Contact: Joe Administrator, x8888

Choose the operation you want to perform:
- SETUP - Configure SNMP.
[]>

example.com>
```

CHAPTER 17

Reporting

This chapter contains the following information:

REPORTING OVERVIEW

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. For more information see, “Exporting Report Data” on page 329.

SCHEDULING REPORTS

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, or previous month. 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).

Adding a Scheduled Report

Use the Monitor > Reports > Add Scheduled Report page to schedule reporting for Anti-Malware, Client Malware Risk, Client Web Activity, L4 Traffic Monitor, Overview, URL Categories, Web Reputation Filters, Web Site Activity.

Figure 17-1 Scheduling Reports

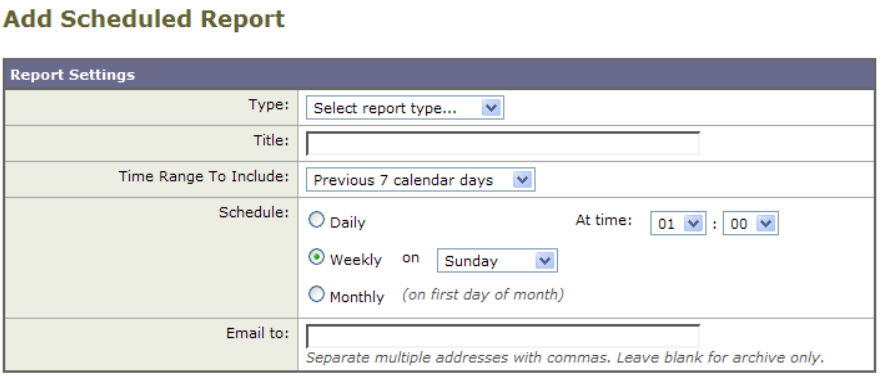

To create a scheduled report:

1. Select a report type.

2.Enter a title for the report. To avoid creating multiple reports with the same name, consider using a descriptive title.

3.Select a time range for the data included in the report.

4.Specify report options, if available. Some reports do not have report options.

5.Specify scheduling and delivery options. If you do not specify an email address, the report is archived only.

6.Submit and commit your changes.

Editing Scheduled Reports

To edit reports, select the report title from the list on the Monitor > Report Scheduling page, modify settings then submit and commit your changes.

Deleting Scheduled Reports

To delete reports, go to the Monitor > Report Scheduling 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.

ON-DEMAND REPORTS

The Generate Report Now option on the Monitor > Archived Reports page allows you to generate on-demand data displays for each report type. To generate a report:

1. Select **Generate Report Now**

   Figure 17-2 Generating an On-Demand Report

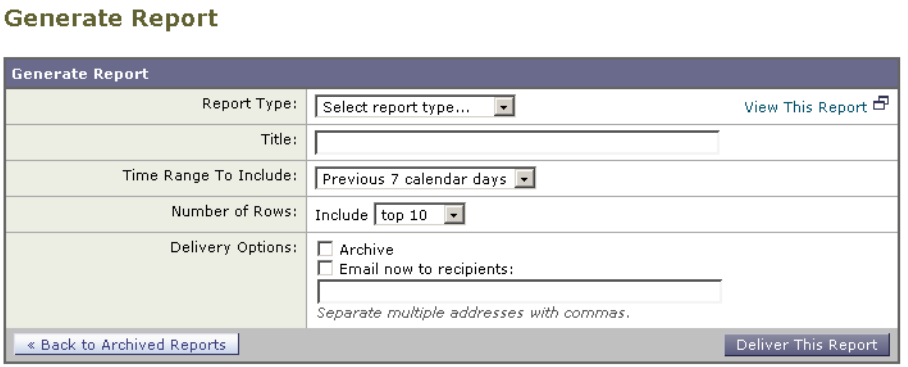

2. Select a report type and edit the title, if necessary. To avoid creating multiple reports with the same name, consider using a descriptive title.
3. Select a time range for the data included in the report.
4. Specify report options, if available.
5. Select whether to archive the report (if so, the report will appear on the Archived Reports page).
6. Specify whether to email the report, and list the email addresses of the recipients.
7. **View This Report** immediately displays the information without having it archived or forwarded to an email distribution list.
8. **Deliver this Report** generates the report.
9. **Commit** your changes.

ARCHIVING REPORTS

The Monitor > Archived Reports page lists available archived reports. Report names in the Report Title column are interactive and link to a view of each report. The Show menu filters the types of reports that are listed. Additionally, interactive column headings can be used to sort the data in each 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.

EXPORTING REPORT DATA

Export links on the display pages will export raw data to a comma-separated values (CSV) file, that you can access and manipulate using database applications such as, Microsoft Excel.

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 17-1 Viewing Raw Data Entries

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

Note — Category headers are different for each type of report.

CHAPTER 18

Logging

This chapter contains the following information:

LOGGING OVERVIEW

You can use log files to monitor web traffic. To configure the appliance to create log files, you create log subscriptions. A log subscription is an appliance configuration that associates a log file type with a name, logging level, and other parameters, such as size and destination information. You can subscribe to a variety of log file types. For more information about log subscriptions, see "Working with Log Subscriptions" on page 336.

In typical appliance monitoring, the appliance administrator usually reads the following log files:

- **Access log.** Records all Web Proxy filtering and scanning activity. For more information about the access log, see "Access Log File" on page 343.
- **Traffic Monitor log.** Records all L4 Traffic Monitor activity. For more information about the traffic monitor log, see "Traffic Monitor Log" on page 355.

The appliance also creates other log file types, such as the system log file. You might want to read other log files to troubleshoot appliance errors. For a list of each type, see "Log File Types" on page 332.

The appliance provides several options for customizing the type of information recorded in the access log. For more information, see "Custom Formatting" on page 356.

Log File Types

The log file type indicates what information is recorded in the generated log, such as web traffic or system data. By default, the Web Security appliance has log subscriptions for most log file types already created. However, there are some log file types that specific to troubleshooting the Web Proxy. Those logs are not created by default. For more information on those log file types, see "Web Proxy Logging" on page 334.

Table 18-1 lists the Web Security appliance log file types created by default.

Table 18-1 Default Log File Types

Log File Type	Description
Access Logs	Records web proxy client history.
Authentication Framework Logs	Records authentication history and messages.
CLI Audit Logs	Records a historical audit of command line interface activity.
Default Proxy Logs	Records errors related to the Web Proxy. This is the most basic of all Web Proxy related logs. To troubleshoot more specific aspects related to the Web Proxy, create a log subscription for the applicable Web Proxy module. For more information about Web Proxy logging, see "Web Proxy Logging" on page 334.

Table 18-1 Default Log File Types (Continued)

Log File Type	Description
External Authentication Logs	Records messages related to using the external authentication feature, such as communication success or failure with the external authentication server. Even with external authentication is disabled, this log contains messages about local users successfully or failing logging in. For more information on external authentication, see "Using External Authentication" on page 374.
Feedback Logs	Records the web users reporting misclassified pages.
GUI Logs (Graphical User Interface)	Records history of page refreshes in the web interface.
Logging Logs	Records errors related to log management.
McAfee Logs	Records the status of anti-malware scanning activity from the McAfee scanning engine.
NTP Logs (Network Time Protocol)	Records changes to the system time made by the Network Time Protocol.
PAC File Hosting Daemon Logs	Records proxy auto-config (PAC) file usage by clients.
Proxy Bypass Logs	Records transactions that bypass the Web Proxy.
Reporting Logs	Records a history of report generation.
Reporting Query Logs	Records errors related to report generation.
SHD Logs (System Health Daemon)	Records a history of the health of system services and a history of unexpected daemon restarts.
Status Logs	Records information related to the system, such as feature key downloads.
System Logs	Records DNS, error, and commit activity.
Traffic Monitor Error Logs	Records L4TM interface and capture errors.
Traffic Monitor Logs	Records sites added to the L4TM block and allow lists.
Updater Logs	Records a history of WBRS and other updates.
WBNP Logs (SenderBase Network Participation)	Records a history of SenderBase network participation uploads to the SenderBase network.

Table 18-1 Default Log File Types (Continued)

Log File Type	Description
WBRS Logs (Web Reputation Score)	Records the status of Web Reputation Filters service, such as whether or not the service is running.
Web Categorization Logs	Records the status of the IronPort URL Filters service, such as whether or not the service is running.
Webroot Logs	Records the status of anti-malware scanning activity from the Webroot scanning engine.
Welcome Page Acknowledgement Logs	Records a history of web clients who click the Accept button on the end-user acknowledgement page.

Web Proxy Logging

By default, the Web Security appliance has one log subscription created for Web Proxy logging messages, the "Default Proxy Logs." The Web Proxy information stored in this log covers all aspects, or modules, of the Web Proxy. The appliance also includes log file types for each Web Proxy module so you can read more specific debug information for each module without cluttering up the Default Proxy Logs.

If a user or administrator encounters an issue with the Web Proxy behavior, read the Default Proxy Logs first. If you see a log entry that you suspect might be the symptom of an issue, then you can create a log subscription for the relevant specific Web Proxy module. Then read that proxy log to help troubleshoot the problem.

You can create log subscriptions of these proxy module logs in web interface or in the CLI. However, you can only create the Request Debug Logs in the CLI.

Table 18-2 describes the Web Proxy module log types.

Table 18-2 Web Proxy Module Log File Types

Web Proxy Module Log File Type	Description
Access Control Engine Logs	Records messages related to the Web Proxy ACL (access control list) evaluation engine.
Configuration Logs	Records messages related to the Web Proxy configuration management system.
Connection Management Logs	Records messages related to the Web Proxy connection management system.
Disk Manager Logs	Records Web Proxy messages related to writing to the cache on disk.

Table 18-2 Web Proxy Module Log File Types (Continued)

Web Proxy Module Log File Type	Description
HTTPS Logs	Records Web Proxy messages specific to the HTTPS proxy (when HTTPS scanning is enabled).
License Module Logs	Records messages related to the Web Proxy's license and feature key handling system.
Logging Framework Logs	Records messages related to the Web Proxy's logging system.
McAfee Integration Framework Logs	Records messages related to communication between the Web Proxy and the McAfee scanning engine.
Memory Manager Logs	Records Web Proxy messages related to managing all memory including the in-memory cache for the Web Proxy process.
Miscellaneous Proxy Modules Logs	Records Web Proxy messages that are mostly used by developers or customer support.
Request Debug Logs	Records very detailed debug information on a specific HTTP transaction from all Web Proxy module log types. You might want to create this log subscription to troubleshoot a proxy issue with a particular transaction without creating all other proxy log subscriptions. **Note:** You can create this log subscription in the CLI only.
SNMP Module Logs	Records Web Proxy messages related to interacting with the SNMP monitoring system.
WBRS Framework Logs (Web Reputation Score)	Records messages related to communication between the Web Proxy and the Web Reputation Filters.
WCCP Module Logs	Records Web Proxy messages related to implementing WCCP.
Webcat Integration Framework Logs	Records messages related to communication between the Web Proxy and the IronPort URL Filters.
Webroot Integration Framework Logs	Records messages related to communication between the Web Proxy and the Webroot scanning engine.

WORKING WITH LOG SUBSCRIPTIONS

A log subscription is an appliance configuration that specifies the type of log file to create and other factors, such as the log file name and method of retrieving the log file. Use the System Administration > Log Subscriptions page to configure log file subscriptions.

Figure 18-1 shows the Log Subscriptions page where you work with log subscriptions.

Figure 18-1 Log File Subscriptions

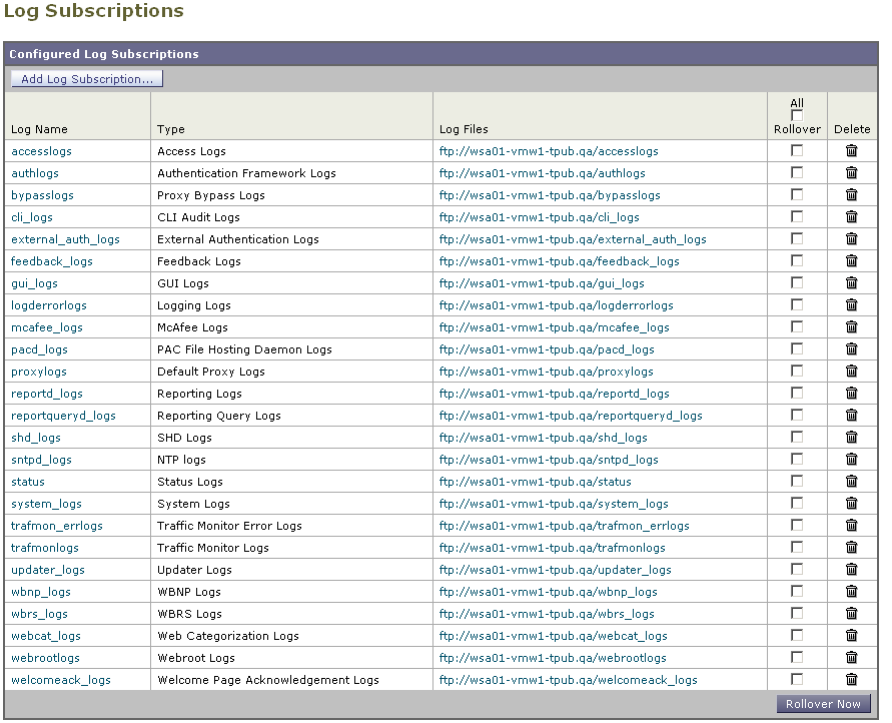

By default, the appliance is configured with one log subscription for every log type. You can add, edit, or delete log subscriptions. You can retrieve log files from the appliance using SCP, FTP, or Syslog.

You can create multiple log subscriptions for each type of log file except the access log. You can only create one access log subscription.

The appliance includes more options when configuring the access log:

- **Include additional information in each log entry.** For more information about customizing the access log, see “Custom Formatting” on page 356.

- **Choose the format of the information.** You can choose among the following format options:
  - Apache
  - Squid
  - Squid Details

Log File Name and Appliance Directory Structure

The appliance creates a directory for each log subscription based on the log subscription name. The name of the log file in the directory is composed of the following information:

- Log file name specified in the log subscription
- Timestamp when the log file was started
- A single-character status code, either `.c` (signifying current) or `.s` (signifying saved)

The filename of logs are made using the following formula:

```
/LogSubscriptionName/LogFilename.@timestamp.statuscode
```

Note — You should only transfer log files with the saved status.

Rolling Over Log Subscriptions

AsyncOS rolls over log subscriptions based on settings you make in each log subscription. Rolling over a log subscription is an AsyncOS process that accomplishes the following tasks:

- 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 when the log retrieval method is push-based. For a list of the log retrieval methods, see Table 18-5 on page 340.
- Transfers any previously unsuccessful log files from the same subscription when the log retrieval method is push-based.
- Deletes the oldest file in the log subscription if the total number of files to keep on hand has been exceeded when the log retrieval method is poll-based.

AsyncOS rolls over log subscriptions in the following ways:

- **Manually.** The appliance administrator can manually roll over log subscriptions on demand from either the web interface or the CLI. Use the **Rollover Now** button on the System Administration > Log Subscriptions page, or the `rollovernow` CLI command. The `rollovernow` command allows you to roll over all log files at once or select a specific log file from a list.
- **Automatically.** AsyncOS rolls over log subscriptions based on the first user-specified limit reached: maximum file size or maximum time. Log subscriptions based on the FTP poll retrieval method create files and store them in the FTP directory on the appliance until

they are retrieved from a remote FTP client, or until the system needs to create more space for log files.

To roll over a log subscription in the web interface:

1. Navigate to the System Administration > Log Subscriptions page.
2. Click the check box under the Rollover column for each log subscription you want to roll over.
3. Click **Rollover Now**.

Viewing the Most Recent Log Files

You can view a the most recent version of a log file from the following locations:

- **Web interface.** On the System Administration > Log Subscriptions page, click the name of the log subscription in the Log Files column of the list of log subscriptions. When you click the link to the log subscription, AsyncOS prompts you to enter your password. Then it lists the available log files for that subscription. Click one of the log files to view it in your browser or to save it to disk.
- **Command line interface.** Use the `tail` CLI command. AsyncOS displays the configured log subscriptions and prompts you to select the log subscription to view. Use Ctrl+C to exit from the `tail` command at any time.

Configuring Host Keys

Use the `logconfig -> hostkeyconfig` subcommand to manage host keys for use with SSH when pushing log files to other servers from the Web Security appliance. SSH servers must have a pair of host keys, one private and one public. The private host key resides on the SSH server and cannot be read by remote machines. The public host key is distributed to any client machine that needs to interact with the SSH server.

The `hostkeyconfig` subcommand performs the following functions:

Table 18-3 Managing Host Keys—List of Subcommands

Command	Description
New	Add a new key.
Scan	Automatically download a host 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	Displays the public key of the system account that pushes the logs to the remote machine. This is the same key that is displayed when setting up an SCP push subscription. This is the value to place in the remote system's 'authorized_keys' file.

Adding and Editing Log Subscriptions

To add or edit a log subscription:

1. Navigate to the System Administration > Log Subscriptions page.
2. To add a log subscription, click **Add Log Subscription**. Or, to edit a log subscription, click the name of the log file in the Log Name field.

   The New Log Subscription page or Edit Log Subscription page appears.
3. Select the type of log to associate with this subscription from the Log Type field.
4. Enter a name for the log subscription in the Log Name field.

   The appliance uses this name for the directory on the appliance that will contain the log file.
5. If you are creating an access log, configure the following options:

Access Log Option	Description
Log Style	Choose the log format to use, either Squid, Apache, or Squid Details.
Custom Fields	Optionally, enter the other type of information to include in each access log entry. For more information, see "Custom Formatting" on page 356.

6. Enter a name for the log file in the File Name field.
7. Enter the maximum file size in bytes the log file can be in the Maximum File Size field. After the number, enter "M" to specify Megabytes or "K" for Kilobytes.
8. Choose the amount of detail to include in the log file in the Log Level field.

   More detailed settings create larger log files and have a greater impact on system performance. More detailed settings include all the messages contained in less detailed settings, plus additional messages. As the level of detail increases, system performance decreases.

   Table 18-4 describes the levels of detail you can choose in the Log Level field.

Table 18-4 Logging Levels

Log Level	Description
Critical	This is the least detailed setting. This level only includes errors. Using this setting will not allow you to monitor performance and other important activities. However, the log files will not reach their maximum size as quickly. This log level is equivalent to the syslog level "Alert."

Table 18-4 Logging Levels (Continued)

Log Level	Description
Warning	This level includes all errors and warnings created by the system. Using this setting will not allow you to monitor performance and other important activities. This log level is equivalent to the syslog level "Warning."
Information	This level includes the detailed system operations. This is the default. This log level is equivalent to the syslog level "Info."
Debug	This level includes data useful for debugging system problems. Use the Debug log level when you are trying to discover the cause of an error. Use this setting temporarily, and then return to the default level. This log level is equivalent to the syslog level "Debug."
Trace	This is the most detailed setting. This level includes a complete record of system operations and activity. The Trace log level is recommended only for developers. Using this level causes a serious degradation of system performance and is not recommended. This log level is equivalent to the syslog level "Debug."

9. Choose how to retrieve the log file from the appliance in the Retrieval Method field.

   Table 18-5 describes the different ways you can retrieve log files:

Table 18-5 Log Transfer Protocols

Retrieval Method	Description
FTP on Appliance (FTP Poll)	This method requires a remote FTP client accessing the appliance to retrieve log files using an admin or operator user's username and password. When you choose this method, you must enter the maximum number of log files to store on the appliance. When the maximum number is reached, the system deletes the oldest file. This is the default.

Table 18-5 Log Transfer Protocols (Continued)

Retrieval Method	Description
FTP on Remote Server (FTP Push)	This method periodically pushes log files to an FTP server on a remote computer. When you choose this method, you must enter the following information: • Maximum time between file transfers • FTP server host name • Directory on FTP server to store the log file • Username and password of a user that has permission to connect to the FTP server **Note:** AsyncOS for Web only supports passive mode for remote FTP servers. It cannot push log files to an FTP server in active mode.
SCP on Remote Server (SCP Push)	This method periodically pushes log files using the secure copy protocol to an SCP server on a remote computer. This method requires an SSH SCP server on a remote computer using the SSH1 or SSH2 protocol. When you choose this method, you must enter the following information: • Maximum time between file transfers • Protocol to use for transmission, either SSH1 or SSH2 • SCP server host name • Directory on SCP server to store the log file • Username and password of a user that has permission to connect to the SCP server Choose whether or not to enable host key checking.
Syslog Push	This method sends log messages to a remote syslog server. This method conforms to RFC 3164. The appliance uses port 514. When you choose this method, you must enter the following information: • Syslog server host name • Protocol to use for transmission, either UDP or TCP • Facility to use with the log You can only choose syslog for text-based logs.

10. Submit and commit your changes.
11. If you chose SCP as the retrieval method, the appliance displays an SSH key to you must place on the SCP server host.

Deleting a Log Subscription

To delete a log subscription:

1. Navigate to the System Administration > Log Subscriptions page.
2. Click the icon under the Delete column for the log subscription you want to delete.
3. Submit and commit your changes.

ACCESS LOG FILE

The access log file provides a descriptive record of all Web Proxy filtering and scanning activity. Access log file entries display a record of how the appliance handled each transaction. You can view the access log file from the System Administration > Log Subscriptions page.

The following text is an example access log file entry for a single transaction:

```
1149143109.100 97 172.xx.xx.xx TCP_MISS/200 8187 GET http://
my.site.com/ - DIRECT/my.site.com text/plain ALLOW_WBRS-
AccessPolicyGroup-IdentityPolicyGroup-RoutingPolicyGroup <CTGY,9.9,-
,-,-,-,-,-,-,-,-,->
```

Table 18-6 describes the different fields in the access log file entry.

Table 18-6 Access Log File Entry

Field Value	Field Description
`1149143109.100`	Timestamp since UNIX epoch.
`97`	Elapsed time (latency) in milliseconds.
`172.xx.xx.xx`	Client IP address.
`TCP_MISS/`	Transaction result code. For more information, see "Transaction Result Codes" on page 344.
`200`	HTTP response code.
`8187`	Response size (headers + body).
`GET http://my.website.com/`	First line of the request.
`-`	Authenticated username.
`DIRECT/my.website.com`	Code that describes from which server was contacted for the retrieving the request content. Most common values include: • **NONE.** The Web Proxy had the content, so it did not contact any other server to retrieve the content. • **DIRECT.** The Web Proxy went to the server named in the request to get the content. • **DEFAULT_PARENT.** The Web Proxy went to its primary parent proxy to get the content.
`text/plain`	Response body MIME type.

Table 18-6 Access Log File Entry (Continued)

Field Value	Field Description
ALLOW_WBRS	ACL decision tag. For more information, see "ACL Decision Tags" on page 345.
AccessPolicyGroup	Access policy group name. When the transaction matches the global access policy, this value is "DefaultGroup." Any space in the policy group name is replaced with an underscore (_).
IdentityPolicyGroup	Identity policy group name. This value is omitted when the transaction matches a policy group that uses "All Identities," including the Global Identity Policy. Any space in the policy group name is replaced with an underscore (_).
RoutingPolicyGroup	Routing policy group name as *ProxyGroupName/ProxyServerName*. When the transaction matches the global routing policy, this value is "DefaultRouting." When no upstream proxy server is used, this value is "DIRECT." Any space in the policy group name is replaced with an underscore (_).
<CTGY,9.9,-,-,-,-,-,-,-,-,-,-,->	Web Reputation filtering and Anti-Malware scanning information. For more information about the values included within the angled brackets, see "Understanding Web Reputation and Anti-Malware Information" on page 346.

Transaction Result Codes

Transaction result codes in the access log file describe how the appliance resolves client requests. For example, if a request for an object can be resolved from the cache, the result code is TCP_HIT. However, if the object is not in the cache and the appliance pulls the object from an origin server, the result code is TCP_MISS. The following table describes transaction result codes.

Table 18-7 Transaction Result Codes

Result Code	Description
TCP_HIT	The object requested was fetched from the disk cache.

Table 18-7 Transaction Result Codes (Continued)

Result Code	Description
TCP_IMS_HIT	The client sent an IMS (If-Modified-Since) request for an object and the object was found in the cache. The proxy responds with a 304 response.
TCP_MEM_HIT	The object requested was fetched from the memory cache.
TCP_MISS	The object was not found in the cache, so it was fetched from the origin server.
TCP_REFRESH_HIT	The object was in the cache, but had expired. The proxy sent an IMS (If-Modified-Since) request to the origin server, and the server confirmed that the object has not been modified. Therefore, the appliance fetched the object from either the disk or memory cache.
TCP_CLIENT_REFRESH_MISS	The client sent a "don't fetch response from cache" request by issuing the 'Pragma: no-cache' header. Due to this header from the client, the appliance fetched the object from the origin server.
TCP_DENIED	The client request was denied due to access policies.
NONE	There was an error in the transaction. For example, a DNS failure or gateway timeout.

ACL Decision Tags

An ACL decision tag is a field in an access log entry that indicates how the DVS engine handled the transaction. It includes information from the Web Reputation filters, URL categories, and the scanning engines.

Table 18-8 describes the ACL decision tag values.

Table 18-8 ACL Decision Tag Values

ACL Decision Tag	Description
ALLOW_ADMIN	The DVS engine allowed the transaction based on Applications settings for the access policy group.
BLOCK_ADMIN	The DVS engine blocked the transaction based on Applications or Objects settings for the access policy group.
BLOCK_WEBCAT	The DVS engine blocked the transaction based on URL category filtering settings for the access policy group.
ALLOW_WBRS	The DVS engine allowed the transaction based on the Web Reputation filter settings for the access policy group.

Table 18-8 ACL Decision Tag Values (Continued)

ACL Decision Tag	Description
BLOCK_WBRS	The DVS engine blocked the transaction based on the Web Reputation filter settings for the access policy group.
MONITOR_SUSPECT_USER_AGENT	The DVS engine monitored the transaction based on the Suspect User Agent setting for the access policy group.
BLOCK_SUSPECT_USER_AGENT	The DVS engine blocked the transaction based on the Suspect User Agent setting for the access policy group.
MONITOR_AMW_REQ	The DVS engine suspects the URL in the HTTP request might not be safe, but it monitored the transaction based on the Anti-Malware settings for the access policy group.
BLOCK_AMW_REQ	The DVS engine suspects the URL in the HTTP request might not be safe, so it blocked the transaction at request time based on the Anti-Malware settings for the access policy group.
MONITOR_AMW_RESP	The DVS engine monitored the server response based on the Anti-Malware settings for the access policy group.
BLOCK_AMW_RESP	The DVS engine blocked the server response based on the Anti-Malware settings for the access policy group.
DEFAULT_CASE	The DVS engine allowed the client to access the server because none of the AsyncOS services, such as Web Reputation or anti-malware scanning, took any action on the transaction.
OTHER	The DVS engine did not complete the request due to an error, such as an authorization failure, server disconnect, or an abort from the client.
NO_PASSWORD	The user failed authentication.

Understanding Web Reputation and Anti-Malware Information

The access log file entries aggregate and display the results of Web Reputation filtering and anti-malware scanning. The appliance displays this information in angled brackets at the end of each access log entry.

The following text is the Web Reputation filtering and anti-malware scanning information from an access log file entry. In this example, the Webroot scanning engine found the malware:

```
<nc,ns,13,"Comedy-Planet",-,2116,363786,-,-,-,-,->
```

The following text is the Web Reputation filtering and anti-malware scanning information from an access log file entry. In this example, the McAfee scanning engine found the malware:

```
<nc,ns,0,-,-,-,-,23,"CP22.EXE",0,1,1,"Generic Downloader.ab">
```

Note — For an example of a whole access log file entry, see "Access Log File" on page 343.

Table 18-9 describes the different fields in the Web Reputation filtering and anti-malware scanning section of each access log file entry.

Table 18-9 Access Log File Entry — Web Reputation and Anti-Malware Information

Field Value Example 1	Field Value Example 2	Description
nc	nc	Abbreviated URL category name. For a list of URL category abbreviations, see "URL Category Abbreviations" on page 350.
ns	ns	Web Reputation filters score. This field either shows the score as a number, "ns" for "no score," or "dns" when there is a DNS lookup error.
13	0	The malware scanning verdict Webroot passed to the DVS engine. Applies to responses detected by Webroot only. For more information, see "Malware Scanning Verdict Values" on page 353.
Comedy-Planet	-	Name of the spyware that is associated with the object. Applies to responses detected by Webroot only.
-	-	The Webroot specific value associated with the Threat Risk Threshold (TRT) value that determines the probability that malware exists. Applies to responses detected by Webroot only.
2166	-	A value that Webroot uses as a threat identifier. IronPort Customer Support may use this value when troubleshooting an issue. Applies to responses detected by Webroot only.
363786	-	A value that Webroot uses as a trace identifier. IronPort Customer Support may use this value when troubleshooting an issue. Applies to responses detected by Webroot only.

Table 18-9 Access Log File Entry — Web Reputation and Anti-Malware Information (Continued)

Field Value Example 1	Field Value Example 2	Description
-	23	The malware scanning verdict McAfee passed to the DVS engine. Applies to responses detected by McAfee only. For more information, see "Malware Scanning Verdict Values" on page 353.
-	CP22.EXE	The name of the file McAfee scanned. Applies to responses detected by McAfee only.
-	0	A value that McAfee uses as a scan error. IronPort Customer Support may use this value when troubleshooting an issue. Applies to responses detected by McAfee only.
-	1	A value that McAfee uses as a detection type. IronPort Customer Support may use this value when troubleshooting an issue. Applies to responses detected by McAfee only.
-	1	A value that McAfee uses as a virus type. IronPort Customer Support may use this value when troubleshooting an issue. Applies to responses detected by McAfee only.
-	Generic Downloader.ab	The name of the virus that McAfee scanned. Applies to responses detected by McAfee only.

Web Reputation Filters Example

In the following example, the URL request was allowed because the URL's Web Reputation score was high enough to qualify to be allowed without being scanned for malware.

```
172.xx.xx.xx TCP_MISS/302 656 GET http://my.website.com/ - DIRECT/
my.website.com text/plain ALLOW_WBRS-PolicyGroup <CTGY,6.0,-,-,-,-,-
,-,-,-,-,-,->
```

In this example, "6.0" is the Web Reputation score. The dash "-" values indicate the request was not forwarded to the DVS engine for anti-malware scanning. The ACL decision tag "ALLOW_WBRS-PolicyGroup" indicates that the request was allowed, and therefore not forwarded for anti-malware scanning, based on this Web Reputation score.

Anti-Malware Request Example

In the following example, the Webroot scanning engine scanned the URL request and assigned a malware scanning verdict based on the URL request. Webroot is the only scanning engine that scans a URL request. For more information about Webroot scanning, see "Webroot Scanning" on page 243.

```
1160078708.895 199 172.xx.xx.xx TCP_DENIED/403 1996 GET http://
www.website.com/path/ - NONE/- - BLOCK_AMW_REQ-PolicyGroup
<nc,ns,10,"Malware",100,-,-,-,-,-,-,-,->
```

In this example, the "nc" stands for "no category" because AsyncOS did not match the URL request to a matching category. The "ns" stands for "no score" because AsyncOS did not find any Web Reputation information about this URL request. Because it did not find any Web Reputation information about the URL, it passed the request to the DVS engine for anti-malware scanning.

The "10" value is the malware scanning verdict that Webroot passes to the DVS engine. ("10" corresponds to generic spyware, as explained in Table 18-11 on page 353.) The "BLOCK_AMW_REQ-PolicyGroup" ACL decision tag shows that Webroot's request-side checking of the URL produced this verdict. The remainder of the fields show the spyware name ("Malware"), threat risk rating ("100"), threat ID ("-"), and trace ID ("-") values, which Webroot derived from its evaluation. In this case, the threat ID and trace ID values are empty ("-") because Webroot did not actually scan a response. All of the McAfee-related values are empty ("-") because the McAfee scanning engine did not scan the URL request.

Anti-Malware Response Example

In the following example, the McAfee scanning engine scanned the server response, assigned a malware scanning verdict based on the server response, and blocked it from the user.

```
1186606394.787 198 172.xx.xx.xx TCP_DENIED/403 1843 GET http://
www.eicar.org/download/eicar.com HTTP/1.1 - NONE/- text/plain
BLOCK_AMW_RESP-PolicyName <Comp,3.0,0,-,-,-,-,27,-,0,1,6,"EICAR test
file">
```

The following list explains the values in this access log entry that show that this transaction was blocked based on the result of the McAfee scanning engine:

- **TCP_DENIED.** The website was denied due to access policies.
- **BLOCK_AMW_RESP-PolicyName.** This transaction matched the "PolicyName" access policy group, and the due to the settings defined in that policy group, the server response was blocked due to detected malware.
- **3.0 in the angled brackets.** The URL received a Web Reputation Score of 3.0, which fell in the score range to scan further.
- **27 in the angled brackets.** The malware scanning verdict McAfee passed to the DVS engine. 27 corresponds to a virus.

- **"EICAR test file".** The name of the virus that McAfee scanned.

URL Category Abbreviations

Table 18-10 lists the abbreviated URL category names that may appear in the in the Web Reputation filtering and anti-malware scanning section of an access log file entry.

Table 18-10 URL Category Abbreviations

URL Category Abbreviation	URL Category	Code
Sear	Search Engines	7503
Spor	Sports	10
Trav	Travel	11
Hobb	Hobbies & Recreation	12
Gamb	Gambling	13
Heal	Health & Medicine	14
News	News	20
Fina	Finance & Investment	25
Fash	Fashion & Beauty	9501
Kids	Kids Sites	7003
Gove	Government	40
Game	Games	1202
Arts	Arts	50
Ente	Entertainment	51
Chat	Chat	3001
Soci	Society & Culture	3003
Job	Job Search & Career Development	60
Reli	Religion	3006
Real	Real Estate	3010
Phil	Philanthropic & Professional Orgs.	9803
Educ	Education	70

Table 18-10 URL Category Abbreviations (Continued)

URL Category Abbreviation	URL Category	Code
Peer	Peer-to-Peer	9801
Infr	Infrastructure	9802
Comp	Computing & Internet	75
Ring	Ringtones/Mobile Phone Downloads	9804
Moto	Motor Vehicles	1101
Poli	Politics	9806
Susp	Suspect/Threat URLs	9101
Hack	Hacking	7504
Sex	Sex Education	1490
Web-	Web-based E-mail	7507
Stre	Streaming Media	7509
Refe	Reference	7001
Adul	Adult/Sexually Explicit	90
Crim	Criminal Activity	91
Into	Intolerance & Hate	92
Viol	Violence	93
Weap	Weapons	94
Inti	Intimate Apparel & Swimwear	95
Pers	Personals & Dating	96
Phot	Photo Searches	97
Prox	Proxies & Translators	98
Host	Hosting Sites	99
Busi	Business	100
Shop	Shopping	80

Table 18-10 URL Category Abbreviations (Continued)

URL Category Abbreviation	URL Category	Code
Food	Food & Dining	3004
Blog	Blogs & Forums	2002
Adve	Advertisements & Popups	76
Down	Downloads	7501
Ille	Illegal Drugs	1403
Alco	Alcohol & Tobacco	1404
Tast	Tasteless & Offensive	9301
-	URL Filtering Bypassed	1073741824
nc	Uncategorized URLs	1073741825
err	URL Filtering Bypassed	1073741826

MALWARE SCANNING VERDICT VALUES

A malware scanning verdict is a value assigned to a URL request or server response that determines the probability that it contains malware. The scanning engines return the malware scanning verdict to the DVS engine so the DVS engine can determine whether to monitor or block the scanned object.

They are the result of proprietary calculations that associate a numerical value to the probability that either the URL request or the response content contains malware. Each malware scanning verdict corresponds to a malware category listed on the "Access Policies: Reputation and Anti-Malware Settings" page when you edit the Web Reputation and Anti-Malware Filtering for a particular access policy.

Both the Webroot and McAfee scanning engines can return malware scanning verdicts to the DVS engine. For more information about how the DVS engine handles malware scanning verdicts, see "IronPort DVS™ (Dynamic Vectoring and Streaming) Engine" on page 240.

Table 18-11 lists the different malware scanning verdict values and each malware category with which they correspond.

Table 18-11 Malware Scanning Verdict Values

Malware Category	Malware Scanning Verdict Value
Unscannable	3
Other Malware	10
Browser Helper Object	12
Adware	13
System Monitor	14
Commercial System Monitor	18
Dialer	19
Hijacker	20
Phishing URL	21
Trojan Downloader	22
Trojan Horse	23
Trojan Phisher	24
Worm	25
Encrypted File	26

Table 18-11 Malware Scanning Verdict Values (Continued)

Malware Category	Malware Scanning Verdict Value
Virus	27

TRAFFIC MONITOR LOG

The L4 Traffic Monitor log file provides a detailed record of monitoring activity. You can view L4 Traffic Monitor log file entries and track updates to firewall block lists and firewall allow lists. Consider the following example log entries:

Example 1

```
172.xx.xx.xx discovered for blocksite.net (blocksite.net) added to
firewall block list.
```

In this example, where a match becomes a block list firewall entry. The L4 Traffic Monitor matched an IP address to a domain name in the block list based on a DNS request which passed through the appliance. The IP address is then entered into the block list for the firewall.

Example 2

```
172.xx.xx.xx discovered for www.allowsite.com (www.allowsite.com) added
to firewall allow list.
```

In this example, a match becomes an allow list firewall entry. The L4 Traffic Monitor matched a domain name entry and added it to the appliance allow list. The IP address is then entered into the allow list for the firewall.

Example 3

```
Firewall noted data from 172.xx.xx.xx to 209.xx.xx.xx
(allowsite.net):80.
```

In this example, the L4 Traffic Monitor logs a record of data that passed between an internal IP address and an external IP address which is on the block list. Also, the L4 Traffic Monitor is set to monitor, not block.

CUSTOM FORMATTING

Using custom format specifiers, you can configure the access log file to capture comprehensive information about web traffic within the network. Format specifiers for the access log file include:

Table 18-12 Custom Format Specifiers

Format Specifier	Description	Enabled by Default
%a	Client IP address	Yes
%A	Authenticated user	Yes
%b	Body bytes returned	No
%c	Response body MIME content type	Yes
%C	content type %C cookie	No
%d	data source (domain)	No
%D	ACL decision tag	Yes
%e	Elapsed time	No
%E	Error type	No
%f	X-Forwarded-For	No
%g	Authorized group names	No
%h	HTTP response code	Yes
%H	Hierarchy retrieval	Yes
%I	Transaction ID	No
%j	Do not cache response code; DCF flags	No
%k	Data source IP address	No
%L	Request local time	No
%M	Cache miss flags	No
%N	Server name or destination hostname	No

Table 18-12 Custom Format Specifiers (Continued)

Format Specifier	Description	Enabled by Default
%p	Destination port number	No
%P	Protocol	No
%r	Request first line - request method, URI, HTTP version	Yes
%R	Referrer	No
%s	Total bytes	Yes
%t	Time stamp	Yes
%T	Timeout	Yes
%u	User agent	No
%w	Action taken For example: TCP_MISS, TCP_HIT	Yes
%W	Result code	No
%x	Latency	Yes
%Xr	Result code	Yes
%Xv	Malware scanning verdict	Yes
%XW	Decoded WBRS score <-10.0-10.0>	Yes
%XC	URL category abbreviation	Yes
%XF	Full name of a custom URL category	No
%<	Request header	No
%>	Response header	No
%:	Event time	No
%:1<	Wait-time for first request byte from new client connection	No
%:h<	Wait-time for complete client header after first byte	No

Table 18-12 Custom Format Specifiers (Continued)

Format Specifier	Description	Enabled by Default
%:b<	Wait-time for complete client body	No
%:<h	Wait-time to write request header to server after first byte	No
%:<b	Wait-time to write request body to server after header	No
%:>1	Wait-time for first response byte from server	No
%:>h	Wait-time for server header after first response byte	No
%:>b	Wait-time for complete response body after header received	No
%:1>	Wait-time for first byte written to client	No
%:h>	Wait-time for complete header written to client	No
%:b>	Wait-time for complete body written to client	No

Configuring Custom Formatting

Use the System Administration > Log Subscriptions page to configure custom formatting for access log file entries. Click the access log file name to edit the access log subscription.

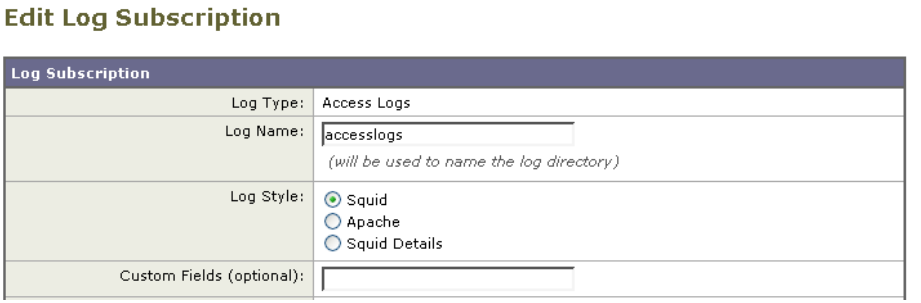

The syntax for entering format specifiers in the Custom Field is as follows:

```
<format_specifier1> <format_specifier2>
```

For example: `%a %b %d`

You can add tokens before the format specifiers to display descriptive text in the access log file. For example:

```
client_IP %a body_bytes %b error_type %E
```

where `client_IP` is the description token for log format specifier `%a`, `body_bytes` is the descriptive token for `%b`, and `error_type` is the descriptive token for `%c`.

CHAPTER 19

System Administration

This chapter contains the following information:

MANAGING THE S-SERIES APPLIANCE

The S-Series appliance provides a variety of tools for managing the system. Functionality on System Administration tab helps you manage the following tasks:

- Appliance configuration
- Feature keys
- Adding, editing, and removing user accounts
- AsyncOS software upgrades
- Updates to security components
- Network settings
- Transparent redirection
- System time

Managing the Appliance Configuration

To archive the current configuration, use the System Administration > Configuration pages to print a summary of appliance settings and create a local copy of the system configuration file. The system configuration file can be used to import a complete configuration or to load a unique sub-section and update specific settings.

Use the System Administration > Configuration File page to load a copy of the current configuration onto the appliance or to copy the configuration to a local host.

To load a copy of the configuration file, paste the configuration directly into the web interface page. At the top of the configuration file you must include the following tag:

```
<?xml version="1.0" encoding="ISO-8859-1"?> <!DOCTYPE config SYSTEM
"config.dtd"> <config> ... your configuration information in valid XML </config>
```

After loading the XML sub-section, submit and commit the update.

Committing Changes to the Appliance Configuration

Each time you modify settings and change appliance behavior using the S-Series web interface, you must first submit your changes and then commit them to the active configuration.

For more information about committing changes, see "Committing and Clearing Changes" on page 10.

SUPPORT COMMANDS

The features in this section are useful when you upgrade the appliance or contact your support provider. You can find the following commands under the Technical Support section of the Support and Help menu:

- **Open a Support Case.** For more information, see "Open a Support Case" on page 363.
- **Remote Access.** For more information, see "Remote Access" on page 364.
- **Packet Capture.** For more information, see "Packet Capture" on page 365.

Open a Support Case

You can use the appliance to send an email to IronPort Customer Support asking for assistance. When the appliance sends the email, it also sends the configuration of the appliance. You can do this in the following ways:

- **CLI.** Use the `supportrequest` command.
- **Web interface.** Use the Support and Help menu > Open a Support Case page.

When you send a support request, you can enter comments describing the issue for which you need support. The appliance must be able to send mail to the Internet to send a support request.

To send a support request in the web interface:

1. From the Support and Help menu, choose Open a Support Case.

Figure 19-1 Open a Technical Support Case Page

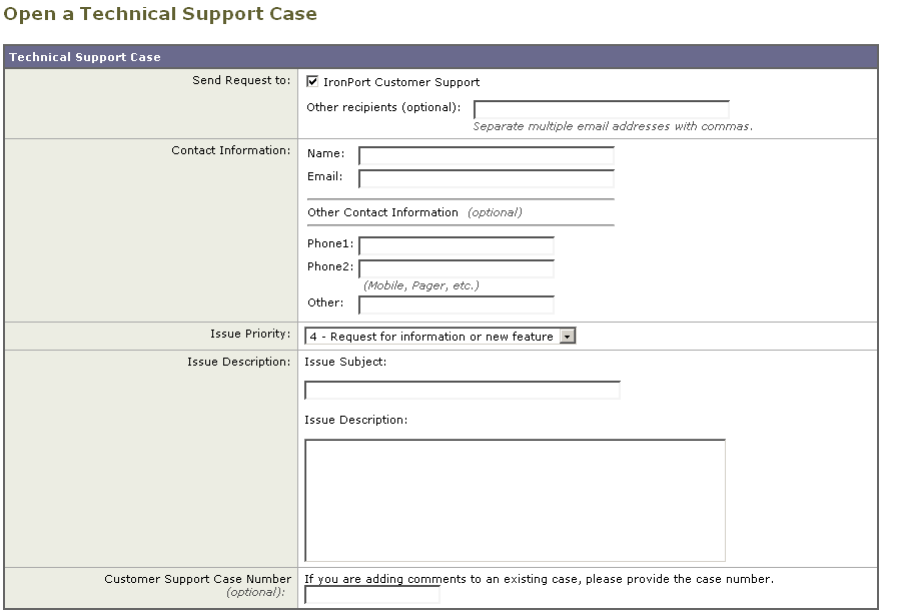

2. In the Other Recipients field, enter other email addresses separated by commas if you want to send this support request to other people.

   By default, the support request (including the configuration file) is sent to IronPort Customer Support (via the checkbox at the top of the form).

3. Enter your contact information, such as name and email.
4. From the Issue Priority field, select the priority of this support request.
5. In the Issue Subject field, enter the text to use in the subject line of the email that will be sent.
6. In the Issue Description field, enter a description of the issue.
7. If you have a customer support ticket already for this issue, enter it.
8. Click **Send**.

   A trouble ticket is automatically created with IronPort. For additional information, see “IronPort Customer Support” on page 10.

Remote Access

Use the Support and Help menu > Remote Access page to allow IronPort Customer Support remote access to the Web Security appliance. Click **Edit Remote Access Settings** to allow IronPort Customer Support to access the appliance.

Figure 19-2 Remote Access Page

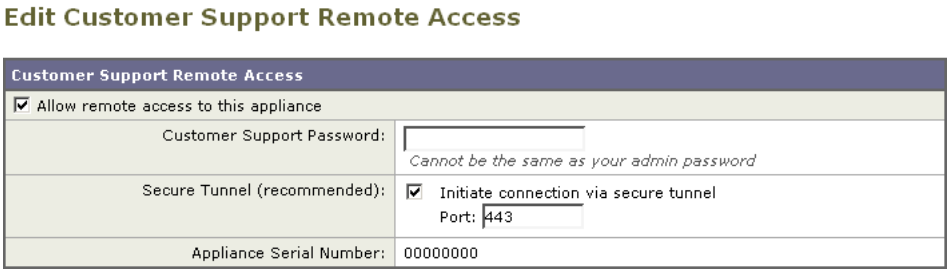

By enabling Remote Access you are activating a special account used by IronPort Customer Support for debugging and general access to the system. This is used by IronPort Customer Support for tasks such as assisting customers in configuring their systems, understanding configurations, and investigating problem reports. You can also use the `techsupport` command in the CLI.

When enabling the "Secure Tunnel," the appliance creates an SSH tunnel over the specified port to the server upgrades.ironport.com. By default this connection is over port 443, which will work in most environments. Once a connection is made to upgrades.ironport.com, IronPort Customer Support is able to use the SSH tunnel to obtain access to the appliance. As long as the connection over port 443 is allowed, this will bypass most firewall restrictions. You can also use the `techsupport tunnel` command in the CLI.

In both the "Remote Access" and "Tunnel" modes, a password is required. It is important to understand that this is *not* the password that will be used to access the system. Once that password and the system serial number are provided to your Customer Support representative, a password used to access the appliance is generated.

Once the techsupport tunnel is enabled, it will remain connected to `upgrades.ironport.com` for 7 days. After 7 days, no new connections can be made using the techsupport tunnel. If there are any existing connections using the tunnel after 7 days, those connections will continue to exist and work. However, once those connections are closed, they will not be able to open again because the techsupport tunnel will have closed after 7 days. The timeout set on the SSH tunnel connection does not apply to the Remote Access account; it will remain active until specifically deactivated.

Packet Capture

Sometimes when you contact IronPort Customer Support with an issue, you may be asked to provide insight into the network activity going into and out of the Web Security appliance. The 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.

Packet Capture

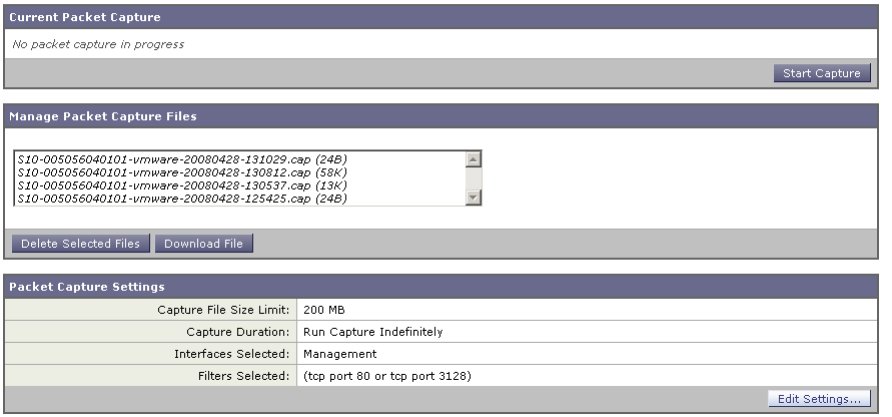

The appliance saves the captured packet activity to a file and stores the file locally. You can configure the maximum packet capture file size, how long to run the packet capture, and on which network interface to run the capture. You can also use a filter to limit the number of packets seen by the packet capture which can make the output more usable on networks with a high volume of traffic. You can send any stored packet capture file using FTP to IronPort Customer Support for debugging and troubleshooting purposes.

The Support and Help > Packet Capture page displays the list of complete packet capture files stored on the hard drive. When a packet capture is running, the web interface shows the status of the capture in progress by showing the current statistics, such as file size and time elapsed.

You can download the packet capture files using the **Download** button in the web interface, or by connecting to the appliance using FTP and retrieving them from the captures directory.

In the CLI, use the `packetcapture` command.

In the web interface, select the Packet Capture option under the Support and Help menu.

Note — The packet capture feature is similar to the Unix tcpdump command.

Starting a Packet Capture

To start a packet capture in the CLI, run the `packetcapture > start` command. If you need to stop a running packet capture, run the `packetcapture > stop` command.

To start a packet capture in the web interface, select the Packet Capture option under the Support and Help menu, and then click **Start Capture**. To stop a running capture, click **Stop Capture**.

Note — The web interface only displays packet captures started in the web interface, not from the CLI. Similarly, the CLI only displays the status of a current packet capture run started in the CLI.

Editing Packet Capture Settings

To edit the packet capture settings in the CLI, run the `packetcapture > setup` command.

To edit packet capture settings in the web interface, select the Packet Capture option under the Support and Help menu, and then click **Edit Settings**.

Table 19-1 describes the packet capture settings you can configure.

Table 19-1 Packet Capture Configuration Options

Option	Description
Capture file size limit	The maximum file size for all packet capture files.
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. **Note:** If the file reaches the maximum size limit before the entire time has elapsed, the existing file is deleted (the data is discarded) and a new file starts with the current packet capture data. • **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 existing file is deleted (the data is discarded) and a new file starts with the current packet capture data. You can always manually stop any packet capture.
Network interface to capture	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 one of the predefined filters to filter by port, source IP address, or destination IP address, or you can create a custom filter using any syntax supported by the Unix tcpdump command.

Note — When you change the packet capture settings without committing the changes and then start a packet capture, AsyncOS uses the new settings. This allows you to use the new settings in the current session without enforcing the settings for future packet capture runs. The settings remain in effect until you clear them.

Figure 19-3 on page 368 shows where you can edit the packet capture settings in the web interface.

Figure 19-3 Editing Packet Capture Settings in the Web Interface

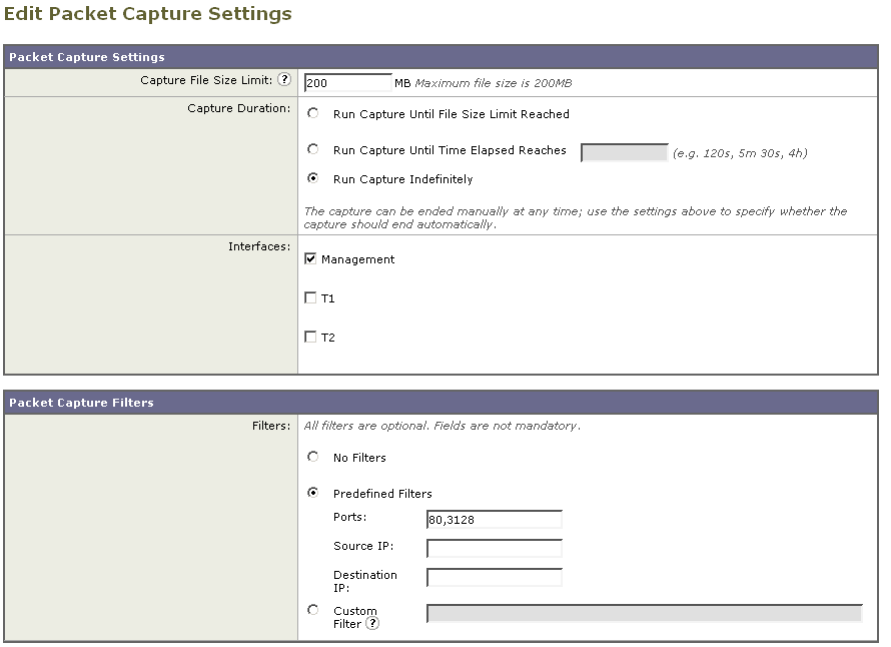

WORKING WITH FEATURE KEYS

Occasionally, your support team may provide a key to enable specific functionality on your system. Use the System Administration > Feature Keys page in the web interface (or the `featurekey` command in the CLI) to enter the key and enable the associated functionality.

Keys are specific to the serial number of your appliance and specific to the feature being enabled (you cannot re-use a key from one system on another system). If you incorrectly enter a key, an error message is generated.

Feature keys functionality is split into two pages: Feature Keys and Feature Key Settings.

Feature Keys Page

The Feature Keys page:

- Lists all active feature keys for the appliance.
- Shows any feature keys that are pending activation.
- Looks for new keys that have been issued (optional, and also can install keys).

A list of the currently enabled features is displayed. The Pending Activation section is a list of feature keys that have been issued for the appliance but have not yet been activated. Your appliance may check periodically for new keys depending on your configuration. You can click **Check for New Keys** to refresh the list of pending keys.

Figure 19-4 The Feature Keys Page

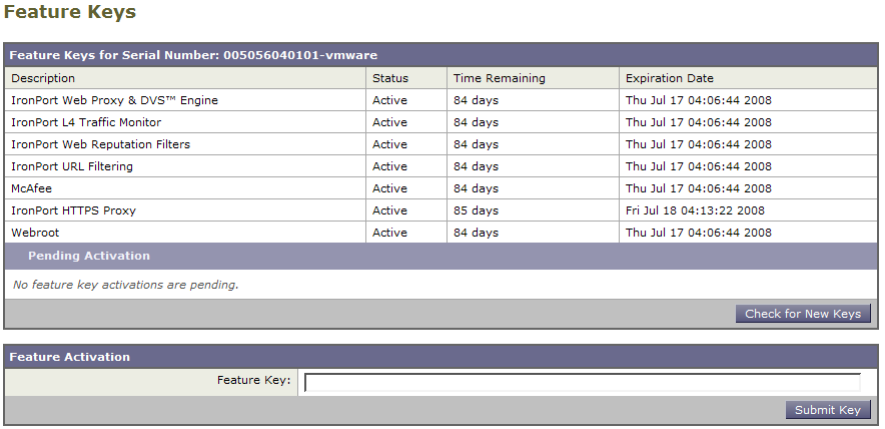

You can also use the `featurekey` CLI command to accomplish the same tasks as on the Feature Keys page.

Feature Key Settings Page

The Feature Key Settings page is used to control whether your appliance checks for and downloads new feature keys, and whether or not those keys are automatically activated.

Figure 19-5 The Feature Key Settings Page

To add a new feature key manually, paste or type the key into the Feature Key field and click **Submit Key**. An error message is displayed if the feature is not added (if the key is incorrect, etc.), otherwise the feature key is added to the display.

To activate a new feature key from the Pending Activation list, select the key (mark the "Select" checkbox) and click **Activate Selected Keys**.

You can configure your IronPort appliance to automatically download and install new keys as they are issued. In this case, the Pending Activation list will always be empty. You can tell AsyncOS to look for new keys at any time by clicking the **Check for New Keys** button, even if you have disabled the automatic checking via the Feature Key Settings page.

You can also use the `featurekeyconfig` CLI command to accomplish the same tasks as on the Feature Key Settings page.

Expired Feature Keys

If the feature key for the feature you are trying to access (via the web interface) has expired, please contact your IronPort representative or support organization.

ADMINISTERING USER ACCOUNTS

The following types of users can log into the Web Security appliance to manage the appliance:

- **Local users.** You can define users locally on the appliance itself. For more information, see "Managing Local Users" on page 371.
- **Users defined in an external system.** You can configure the appliance to connect to an external RADIUS server to authenticate users logging into the appliance. For information, see "Using External Authentication" on page 374.

You can manage local users and connections to external authentication servers using the System Administration > Users page in the web interface, or the `userconfig` command in the CLI.

Figure 19-6 shows where you manage local users and external authentication.

Figure 19-6 System Administration > Users Page

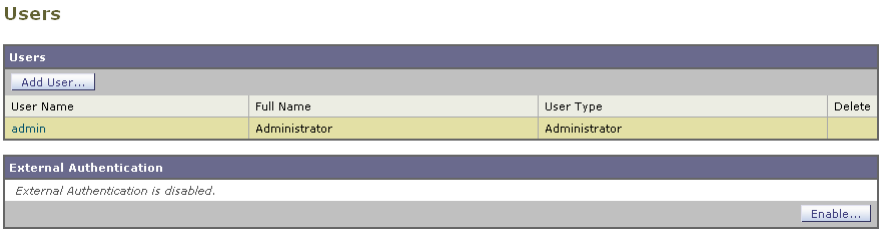

Note — Any user you define can log into the appliance using any method, such as logging into the web interface or using SSH.

Managing Local Users

You can define any number users locally on the Web Security appliance. You can add, edit, and delete local users. Consider the following rules when defining local users:

- User names can contain lowercase letters, numbers, and the dash (-) character.
- User names cannot start with a dash.
- User names cannot be longer than 16 characters.
- Passwords must be at least 6 characters long.
- User names cannot be special names that are reserved by the system, such as "operator" or "root."

The default system admin account has all administrative privileges. You can change the admin account password, but you cannot edit or delete this account.

To create a new user account, specify a user name and a full name, and then assign the user to a group. Each group provides a different level of default permissions. Table 19-2 lists the groups you can assign.

Table 19-2 User Groups

Group	Description
`Administrator`	The administrators group allows full access to all system configuration settings. However, the `upgradecheck` and `upgradeinstall` commands can be issued only from the system defined "admin" account.
`Operator`	The operators group restricts users from creating, editing, or removing user accounts. The operators group also restricts the use of the following commands: • `resetconfig` • `upgradecheck` • `upgradeinstall` • `systemsetup` or running the System Setup Wizard
`Guest`	The guests group users can only view system status information.

After assigning the user to a group, you must specify a password for the new account. Passwords are encrypted when they are reported using the `showconfig` CLI command.

Note — If you have lost the admin user password, contact your support provider.

Adding Local Users

To add a local user:

1. On the System Administration > Users page, click **Add User**.

   The Add Local User page is displayed.

   Figure 19-7 Adding a Local User

   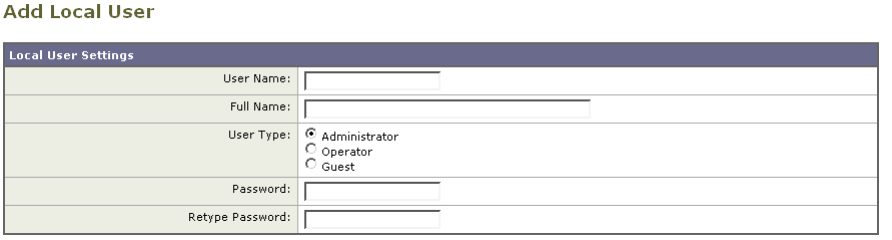

2. Enter a name for the user. Some words are reserved, such as "operator" and "root".
3. Enter a full name for the user.

4. Select a user type. See Table 19-2, "User Groups," on page 372 for more information about user types.
5. Enter a password and retype it.
6. Submit and commit your changes.

Deleting Users

To delete a user:

1. On the System Administration > Users page, click the trash can icon corresponding to the listed user name.
2. Confirm the deletion by clicking **Delete** in the warning dialog that appears.
3. Submit and commit your changes.

Editing Users

To edit a user:

1. On the System Administration > Users page, click the user name.

   The Edit User page is displayed.
2. Make changes to the user.
3. Submit and commit your changes.

Changing Passwords

Users can change their own passwords using the Change Password option under the Options menu located on the top right-hand side of the web interface.

Figure 19-8 shows where you can change the current user password.

Figure 19-8 The Change Password Option

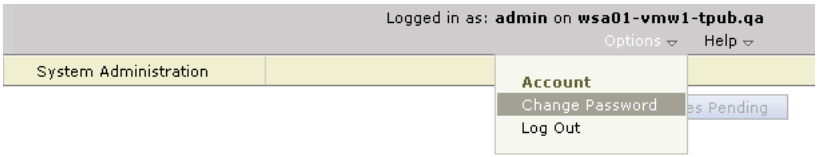

Note — To change the password for the admin account, use the System Administration > Users page or use the `password` or `passwd` command in the CLI. Password changes take effect immediately and do not require a commit.

Monitoring Users from the CLI

The `who`, `whoami`, and `last` commands can be used to monitor user access to the appliance.

- The `who` command lists users, the time of login, idle time, and the remote host from which the user is logged in:

```
example.com> who

Username  Login Time  Idle Time  Remote Host  What
========  ==========  =========  ===========  ====
admin     03:27PM     0s         10.xx.xx.xx  cli
```

- The `whoami` command displays the user name and group information:

```
example.com> whoami

Username: admin
Full Name: Administrator
Groups: admin, operators, config, log, guest
```

- The `last` command displays information about users who have recently logged into the appliance.

```
example.com> last

Username  Remote Host  Login Time        Logout Time       Total Time
========  ===========  ================  ================  ==========
admin     10.xx.xx.xx  Sat May 15 23:42  still logged in   15m
admin     10.xx.xx.xx  Sat May 15 22:52  Sat May 15 23:42  50m
admin     10.xx.xx.xx  Sat May 15 11:02  Sat May 15 14:14  3h 12m
admin     10.xx.xx.xx  Fri May 14 16:29  Fri May 14 17:43  1h 13m
shutdown                                 Fri May 14 16:22
```

Using External Authentication

If you store user information in a RADIUS directory on your network, you can configure the Web Security appliance to use the RADIUS directory to authenticate users logging in to the appliance. You can use external authentication when logging into the appliance using HTTP, HTTPS, SSH, and FTP. To set up the appliance to use an external directory for authentication, use the System Administration > Users page in the web interface or the `userconfig > external` CLI command.

Figure 19-9 shows where you enable external authentication on the System Administration > Users page.

Figure 19-9 Enabling External Authentication

You can configure the appliance to contact multiple external servers for authentication. You might want to define multiple external servers to allow for failover in case one server is temporarily unavailable. When you define multiple external servers, the appliance connects to the servers in the order defined on the appliance.

When external authentication is enabled and a user logs into the Web Security appliance, the appliance first determines if the user is the system defined "admin" account. If not, then the appliance checks the first configured external server to determine if the user is defined there. If the appliance cannot connect to the first external server, the appliance checks the next external server in the list. If the user fails authentication on any external server, the appliance tries to authenticate the user as a local user defined on the Web Security appliance. If the user does not exist on any external server or on the appliance, or if the user enters the wrong password, access to the appliance is denied.

Note — AsyncOS for Web connects to the external server over the M1 interface only.

The Web Security appliance assigns all users in the RADIUS directory to the administrator user group. You cannot assign users to other user groups. When external authentication is enabled and a user successfully authenticates as a local user, the local user has Administrator user group privileges regardless of the configured user type.

To enable external authentication using RADIUS:

1. On the System Administration > Users page, click **Enable**.

   The Edit External Authentication page is displayed.

2. Enable the **Enable External Authentication** option if it is not enabled already.

   Figure 19-10 Enabling External Authentication Using RADIUS

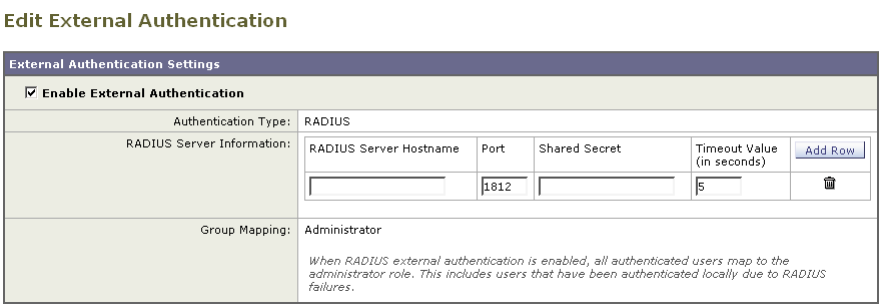

3. Enter the host name for the RADIUS server.
4. Enter the port number for the RADIUS server. The default port number is 1812.
5. Enter the Shared Secret password for the RADIUS server.
6. Enter the number of seconds for the appliance to wait for a response from the server before timing out.

7. Optionally, click **Add Row** to add another RADIUS server. Repeat steps 3-6 for each RADIUS server.

   Note — You can add up to ten RADIUS servers.

8. Submit and commit your changes.

CONFIGURING ADMINISTRATOR SETTINGS

You can configure the Web Security appliance to have stricter access requirements for administrators logging into the appliance. You might want to do this to meet certain organization requirements.

You configure these settings with the `adminaccessconfig` CLI command. You can configure the appliance to:

- Display user-defined text at administrator login.
- Restrict administrator access to certain machines.
- Require stronger SSL ciphers for administrator access.

Configuring Custom Text at Login

Using the `adminaccessconfig > banner` CLI command, you can configure the appliance to display any text you specify when an administrator tries to logs in. You might want to do this to display a banner that informs the user of organizational policies and conditions. The custom banner text appears when an administrator tries to access the appliance through all interfaces, such as the web interface or via FTP.

You can load the custom text by either pasting it into the CLI prompt or by copying it from a file located on the Web Security appliance. To upload the text from a file, you must first transfer the file to the configuration directory on the appliance using FTP.

Configuring IP-Based Administrator Access

Using the `adminaccessconfig > ipaccess` CLI command, you can control from which IP addresses administrators access the Web Security appliance. Administrators can access the appliance from any machine or from machines with an IP address from a list you specify.

When restrict access to an allow list, you can specify IP addresses, subnets, or CIDR addresses.

By default, when you list the addresses that can access the appliance, the IP address of your current machine is listed as the first address in the allow list. You cannot delete the IP address of your current machine from the allow list.

Configuring the SSL Ciphers for Administrator Access

Using the `adminaccessconfig > strictssl` CLI command, you can configure the appliance so administrators log into the web interface on port 8443 using stronger SSL ciphers (greater than 56 bit encryption).

When you configure the appliance to require stronger SSL ciphers, the change only applies to administrators accessing the appliance using HTTPS to manage the appliance. It does not apply to other network traffic connected to the Web Proxy using HTTPS.

CONFIGURING THE RETURN ADDRESS FOR GENERATED MESSAGES

You can configure the return address for mail generated by AsyncOS for 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.

Configure the return address on the System Administration > Return Addresses page.

Figure 19-11 Configuring Return Addresses

To configure the return address for system-generated email messages:

1. Navigate to the System Administration > Return Addresses page.
2. Click **Edit Settings**.

   Figure 19-12 Editing Return Address Settings

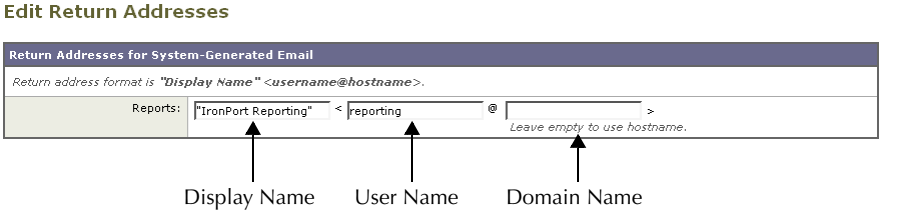

3. For Reports, enter the display name, user name, and domain name in the fields shown in Figure 19-12.
4. Submit and commit your changes.

MANAGING ALERTS

Alerts are email notifications containing information about events occurring on the IronPort appliance. These events can be of varying levels of importance (or severity) from minor (Informational) to major (Critical) and pertain generally to a specific component or feature on the appliance. Alerts are generated by the IronPort appliance. You can specify which alert messages are sent to which users and for which severity of event they are sent. Manage alerts using the System Administration > Alerts page in the web interface or using the `alertconfig` command in the CLI.

Note — To receive alerts and email notifications, you must configure the SMTP relay host that the appliance uses to send the email messages. For information about configuring the SMTP relay host, see "Configuring SMTP Relay Hosts" on page 386.

Alerting Overview

The alerting feature consists of two main parts:

- **Alerts** - consist of an Alert **Recipient** (email addresses for receiving alerts), and the alert notification (severity and alert type) sent to the recipient.
- **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 a specific function (or alert classification) or functions such as a hardware or anti-virus problem, sent to an alert-recipient. An alert recipient is simply an email address to which the alert notifications are sent. The information contained in the notification is determined by an alert classification and a severity. You can specify which alert classifications, at which severity, are sent to any alert recipient. The alerting engine allows for granular control over which alerts are sent to which alert recipients. For example, you can configure the system to send only specific alerts to an alert recipient, configuring an alert recipient to receive notifications only when Critical (severity) information about the System (alert type) is sent. You can also configure general settings (see "Configuring Alert Settings" on page 385).

Alert Classifications

AsyncOS sends the following alert classifications:

Table 19-3 Alert Classifications and Components

Alert Classification	Alert Component
System	System
Hardware	Hardware

Table 19-3 Alert Classifications and Components (Continued)

Alert Classification	Alert Component
Updater	Updater
Web Proxy	Proxy
DVS™ and Anti-Malware	DVS
L4 Traffic Monitor	TrafMon

Severities

Alerts can be sent for the following severities:

- Critical: Requires immediate attention.
- Warning: Problem or error requiring further monitoring and potentially immediate attention.
- Information: Information generated in the routine functioning of this device.

Alert Settings

Alert settings control the general behavior and configuration of alerts, including:

- The RFC 2822 Header From: when sending alerts (enter an address or use the default "alert@<hostname>"). You can also set this via the CLI, using the `alertconfig -> from` command.
- The initial number of seconds to wait before sending a duplicate alert.
- The maximum number of seconds to wait before sending a duplicate alert.
- The status of AutoSupport (enabled or disabled).
- The sending of AutoSupport's weekly status reports to alert recipients set to receive System alerts at the Information level.

Sending Duplicate Alerts

You can specify the initial number of seconds to wait before AsyncOS will send a duplicate alert. If you set this value to 0, duplicate alert summaries are not sent and instead, all duplicate alerts are sent without any delay (this can lead to a large amount of email over a short amount of time). The number of seconds to wait between sending duplicate alerts (alert interval) is increased after each alert is sent. The increase is the number of seconds to wait plus twice the last interval. So a 5 second wait would have alerts sent at 5 seconds, 15, seconds, 35 seconds, 75 seconds, 155 seconds, 315 seconds, etc.

Eventually, the interval could become quite large. You can set a cap on the number of seconds to wait between intervals via the maximum number of seconds to wait before sending a duplicate alert field. For example, if you set the initial value to 5 seconds, and the maximum

value to 60 seconds, alerts would be sent at 5 seconds, 15 seconds, 35 seconds, 60 seconds, 120 seconds, etc.

IronPort AutoSupport

To allow IronPort to better support and design future system changes, the IronPort appliance can be configured to send IronPort Systems a copy of all alert messages generated by the system. This feature, called AutoSupport, is a useful way to allow our team to be proactive in supporting your needs. AutoSupport also sends weekly reports noting the uptime of the system, the output of the `status` command, and the AsyncOS version used.

By default, alert recipients set to receive Information severity level alerts for System alert types will receive a copy of every message sent to 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" on page 385.

Alert Messages

Alert messages are standard email messages. You can configure the Header From: address, but the rest of the message is generated automatically.

Alert From Address

You can configure the Header From: address via the **Edit Settings...** button or via the CLI.

Alert Subject

An alert email message's subject follows this format:

```
Subject: [severity]-[hostname]: ([class]) short message
```

Example Alert Message

```
Date: 23 May 2007 21:10:19 +0000
To: joe@example.com
From: IronPort S650 Alert [alert@example.com]
Subject: Critical <System> example.com: Internal SMTP giving up on
message to jane@company.com with...

The Critical message is:

Internal SMTP giving up on message to jane@company.com with subject
'IronPort Report: Client Web Activity (example.com)': Unrecoverable
error.

Product: IronPort S650 Web Security Appliance
Model: S650
Version: 5.1.0-225
Serial Number: XXXXXXXXXXX-XXXXXXX
Timestamp: Tue May 10 09:39:24 2007

For more information about this error, please see
    http://support.ironport.com
If you desire further information, please contact your support
provider.
```

Managing Alert Recipients

Log in to the S-Series appliance web interface (GUI) and click the System Administration tab. Click the Alerts link in the left menu. For information about how to access the S-Series appliance web interface, see “Accessing the Web Security Appliance” on page 3.

Figure 19-13 The Alerts Page

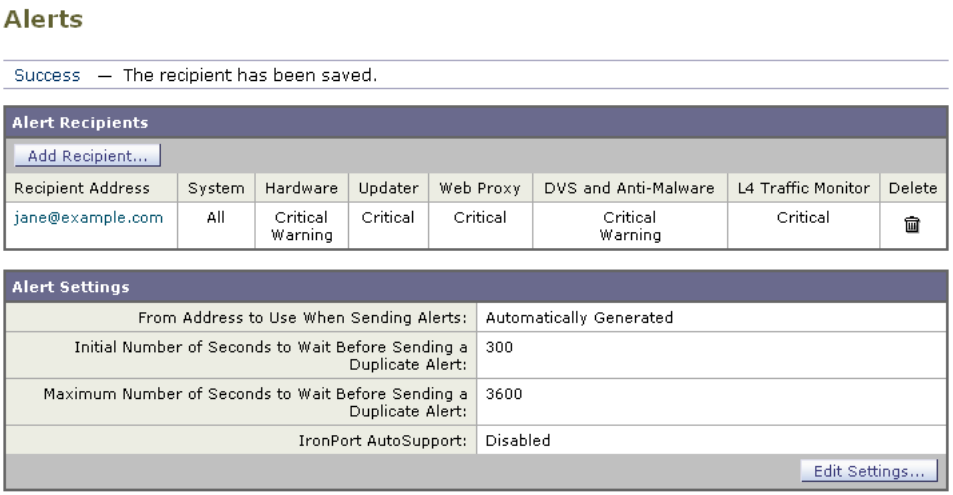

Note — If you enabled AutoSupport during System Setup, the email address you specified will receive alerts for all severities and classes by default. You can change this configuration at any time.

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

To add a new alert recipient:

1. Click **Add Recipient...** on the Alerts page. The Add Alert Recipients page is displayed:

Figure 19-14 Adding a New Alert Recipient

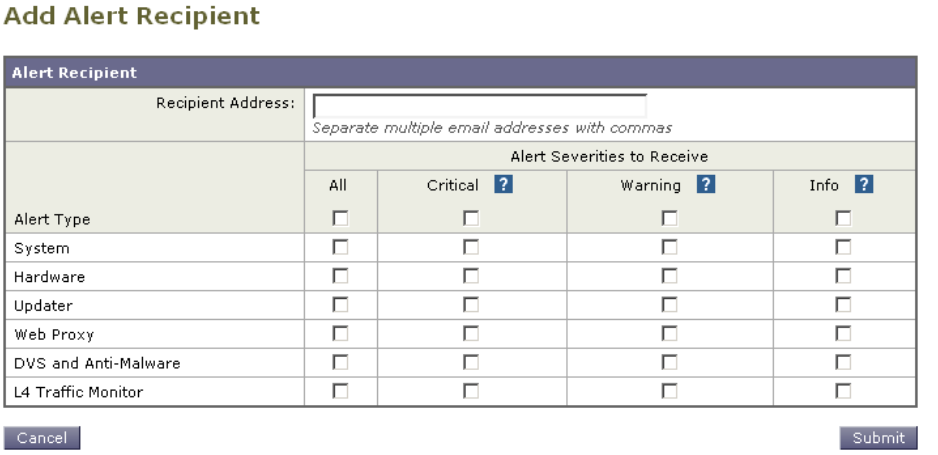

2.Enter the recipient's email address. You can enter multiple addresses, separated by commas.

3.Select which alert severities to receive.

4.Click **Submit** to add the alert recipient.

5.Click the **Commit Changes** button, add an optional comment if necessary, then click **Commit Changes** to save the changes.

Configuring Existing Alert Recipients

To edit an existing alert recipient:

1. Click the alert recipient in the Alert Recipients listing. The Configure Alert Recipient page is displayed.

2.Make changes to the alert recipient.

3.Click **Submit** to edit the alert recipient.

4.Click the **Commit Changes** button, add an optional comment if necessary, then click **Commit Changes** to save the changes.

Deleting Alert Recipients

To delete an alert recipient:

1. Click the trash can icon corresponding to the alert recipient in the Alert Recipient listing.

2.Confirm the deletion by clicking **Delete** in the warning dialog that appears.

3.Click the **Commit Changes** button, add an optional comment if necessary, then click **Commit Changes** to save the changes.

Configuring Alert Settings

Alert settings are global settings, meaning that they affect how all of the alerts behave.

Editing Alert Settings

To edit alert settings:

1. Click **Edit Settings...** on the Alerts page. The Edit Alert Settings page is displayed:

Figure 19-15 Editing Alert Settings

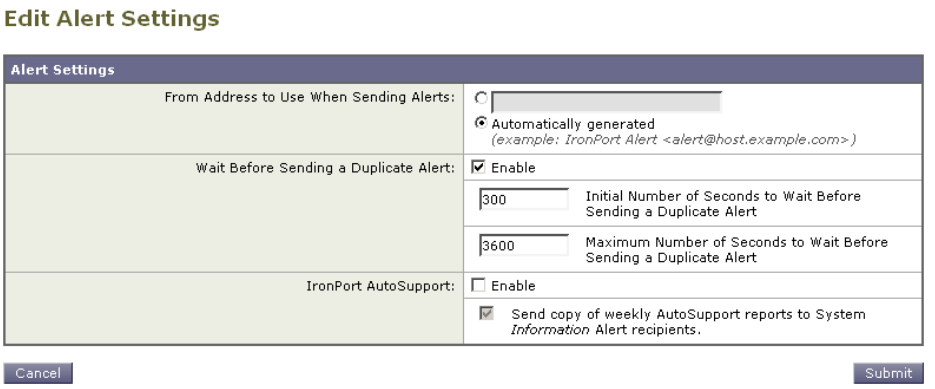

2.Enter a Header From: address to use when sending alerts, or select Automatically Generated (“alert@<hostname>”).

3.Mark the checkbox if you want to specify the number of seconds to wait between sending duplicate alerts. For more information, see “Sending Duplicate Alerts” on page 380.

- 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.

4.You can enable AutoSupport by checking the IronPort AutoSupport option. For more information about AutoSupport, see “IronPort AutoSupport” on page 381.

- If AutoSupport is enabled, the weekly AutoSupport report is sent to alert recipients set to receive System alerts at the Information level. You can disable this via the checkbox.

5.Click **Submit** to edit the alert settings.

6.Click the **Commit Changes** button, add an optional comment if necessary, then click **Commit Changes** to save the changes.

CONFIGURING SMTP RELAY HOSTS

AsyncOS periodically sends system-generated email messages, such as notifications, alerts, and IronPort Customer Support requests. By default, AsyncOS uses information listed in the MX record on your domain to send email. However, if the appliance cannot directly reach the mail servers listed in the MX record, you must configure at least one SMTP relay host on the appliance.

You might want to configure an SMTP relay host in the following scenarios:

- You want the system-generated emails to go to a non-local email address, and port 25 is blocked to outside networks.
- Your mail servers do not allow direct port 25 traffic from internal hosts.

If no SMTP relay host is defined, AsyncOS delivers directly to the mail server for each email address.

Note — If the Web Security appliance cannot communicate with the mail servers listed in the MX record or any of the configured SMTP relay hosts, it cannot send email messages and it writes a message in the log files.

You can configure one or more SMTP relay hosts. You might want to configure multiple SMTP relay hosts for redundancy in case one system becomes unavailable. When you configure multiple SMTP relay hosts, AsyncOS uses the topmost available SMTP relay host. If an SMTP relay host is unavailable, it tries to use the one below it in the list.

You can configure the SMTP relay host from either the web interface or command line interface:

- **Web interface.** Use the Network > Internal SMTP Relay page.
- **Command line interface.** Use the `smtprelay` CLI command.

Configuring SMTP from the Web Interface

Use the Network > Internal SMTP Relay page.

To configure the SMTP relay host from the web interface:

1. Navigate to the Network > Internal SMTP Relay page, and click **Edit Settings**.

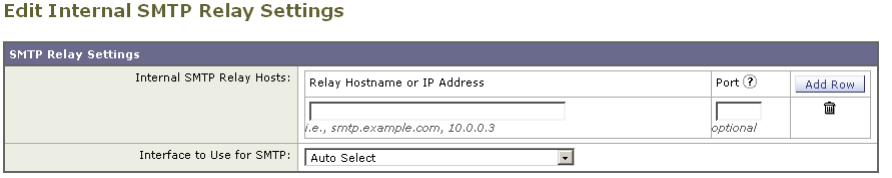

2. Enter the information listed in Table 19-4.

Table 19-4 SMTP Relay Host Settings

Property	Description
Relay Hostname or IP Address	Enter the host name or IP address to use for the SMTP relay
Port	Enter the port for connecting to the SMTP relay. If this property is empty, the appliance uses port 25. This property is optional.
Interface to Use for SMTP	Choose which appliance interface to use for connecting to the SMTP relay. Choose whichever interface is on the same network as the relay system.

3. Optionally, you can add more SMTP relay host information by clicking **Add Row**.
4. Submit and commit your changes.

Configuring SMTP from the CLI

Use the smtprelay command to configure SMTP relay hosts.

For example:

```
example.com> smtprelay

No internal SMTP relay host configured.

Choose the operation you want to perform:

- NEW - Add a new host.

[]> new

Please enter the hostname of your relay host. You may put a colon
after the hostname to indicate a port to use other than 25, such as
“smtp.example.com:547”.
```

UPGRADING THE SYSTEM SOFTWARE

The S-Series appliance supports two methods for upgrading the AsyncOS system software: local and remote.

Remote upgrade: The S-Series appliance downloads the AsyncOS software directly from the IronPort update servers.

Local upgrade: The S-Series appliance downloads the AsyncOS software from the IronPort update server to a server in your own network, and then upgrades the appliance from the local host.

Use the System Administration > Upgrades page in the web interface to configure an upgrade method or use the `upgradeconfig` command from the CLI.

Note — After performing an upgrade, use the System Administration > Configuration File pages or the `saveconfig` command to archive the current configuration file.

Upgrading from a Remote Host

To upgrade the AsyncOS software from a remote location:

1. Obtain a specific URL for a downloadable AsyncOS update image from your support provider.
2. Use a web browser to download the upgrade image. For example:

   ```
   http://downloads.ironport.com/asyncos/upgrade/?serial=serial_number
   ```

   where `serial_number` is the serial number of your IronPort appliance.

Upgrading from a Local Server

To download the AsyncOS software to an internal server and then upgrade the appliance from within your own network:

1. Configure a local server to retrieve the software and host the update. The IronPort update servers are located at the following URL:

   ```
   http://downloads.ironport.com/
   ```

2. Download the upgrade image.
3. Upgrade the appliance using the System Administration > Upgrades page, or from the CLI:
   - Run the `upgradeconfig` command
   - Run the `upgrade` command

Configuring Upgrades

To configure updates using the web interface, use the System Administration > Upgrade Settings page and click **Edit Upgrade Settings**. Figure 19-16 on page 389 shows the Edit Upgrade Settings page.

Figure 19-16 Configuring Software Upgrades

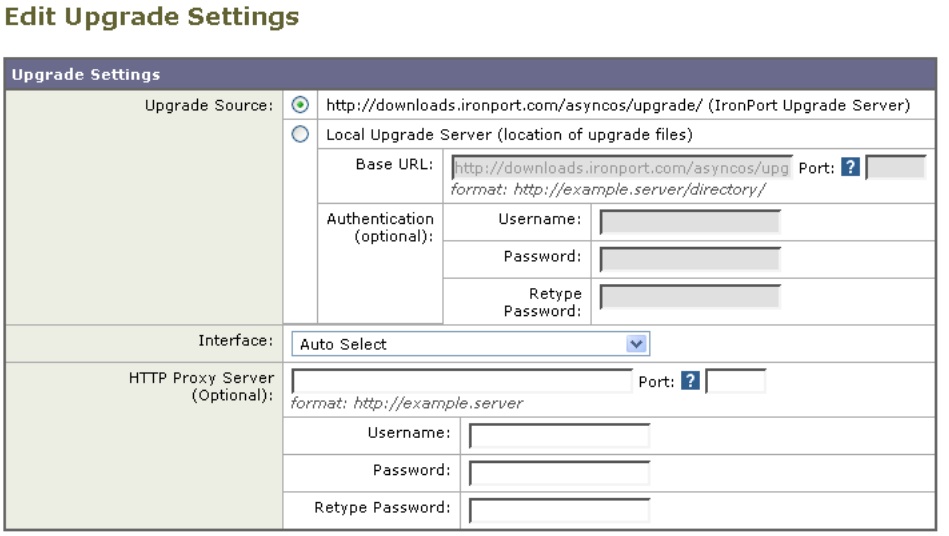

When configuring upgrade settings:

1. Select an upgrade source.
2. For local upgrades, specify the server URL and port number.
3. Select the interface to use for the upgrade.
4. Enter HTTP proxy server information (optional).
5. Submit and commit to save the changes.

Authenticating the Upgrade Server

When configuring the upgrade server, consider how your network configuration can affect the upgrade server. For example, if an upstream proxy is also an authentication server, the appliance might have problems accessing the upgrade list. This issue occurs only when all of the following conditions are true:

- The web proxy is enabled on the P1 or P2 network interface.
- LDAP or NTLM authentication is enabled.
- Upgrades are configured to use the M1 Management Interface.
- Traffic from the M1 interface is routed (WCCP router or switch) through the P1 or P2 interface.

  To resolve this issue:

- Configure the WCCP router or switch to route port 80 connections originating from the Management interface directly to the Internet.
- Avoid configuring an authentication server as an HTTP proxy server for AsyncOS upgrades.
- If LDAP authentication is enabled on the web proxy, explicitly configure HTTP Proxy Server settings on the System Administration > Upgrade Settings > Edit Upgrade Settings page to use the P1 or P2 interface and provide authentication credentials. If the P2 interface is enabled, use P2. If only the P1 interface is enabled, use P1.

Upgrading AsyncOS

To upgrade the AsyncOS system software after configuring your upgrade settings, use the System Administration > System Upgrade > Available Upgrades page.

Upgrading Using the CLI

To upgrade the S-Series appliance using the CLI, run the `upgradeconfig` command to configure an upgrade source. To install an update, run the `upgrade` command.

The upgradeconfig Command

The `upgradeconfig` command specifies a location for AsyncOS updates. When you use the `upgrade` command, the appliance polls IronPort's upgrade servers for the latest update.

```
example.com> upgradeconfig

Choose the operation you want to perform:
- SETUP - Edit upgrade configuration.
[]> setup

Please select the upgrade source you want to use for AsyncOS updates:
1. IronPort upgrade server
2. Local upgrade server
[1]> 2

Please select the location of the upgrade files using the format
(http://optionalname:password@local.server:port/directory/). The
default HTTP port is 80; you do not need to specify the port unless you
wish to use a non-standard port. The optional username/password will
be presented using HTTP BASIC_AUTH.

[]> URL of the local server
```

Note — You can use the `ping` command to ensure that the appliance can contact the local server. You can also use the `telnet` command to telnet to port 80 of the local server to ensure that the local server is listening on that port.

The upgrade Command

The `upgrade` command displays a list of available updates. Select an update from the list:

```
example.com> upgrade

Upgrades available:
1. AsyncOS x.x build xxx upgrade, yyyy-mm-dd
2. AsyncOS x.x build xxx upgrade, yyyy-mm-dd
[2]>

Performing an upgrade will require a reboot of the system after the
upgrade is applied. Do you wish to proceed with the upgrade?  [Y]>

[ ... Possible license agreement statement ... ]

IronPort Web Security Appliance Upgrade

This upgrade will require a reboot of the system after it finishes.
You may log in again after this is done.

Type Return to continue...

Finding partitions... done.
.
.
.

Upgrade done.  It will be in effect after this mandatory reboot.

Rebooting...
Upgrade installation finished.
```

Component Updates

Each security component periodically receives updates to its database tables from the IronPort update server (`https://update-manifests.ironport.com`). Server updates are automated, and the update interval is set by the server as opposed to the appliance. Updates to the database tables occur with a regular degree of frequency and require no administrator intervention.

Manual Updates

Typically, you do not have to perform a manual update to the database tables. In the event a manual update is required, you can modify default settings and configure an update using the options on the System Administration > Component Updates page.

To configure a manual update:

1. Navigate to the System Administration > Component Updates page.

2. Click **Edit Update Settings**.

   The Edit Component Update Settings page appears.

3. Receive the update files from IronPort and install them on a local server.
4. Specify the location of the update files.
5. Initiate the update using the Update Now function key on the component page located on the Security Services tab. For example, Security Services > Web Reputation Filters page.
6. View a record of update activity in the updater log file. Subscribe to the updater log file on the System Administration > Log Subscriptions page.

Note — Updates that are in-progress cannot be interrupted. All in-progress updates must complete before new changes can be applied.

NETWORK SETTINGS

This section describes how to modify the hostname parameter and network settings that were configured during system setup. This section also provides instructions for configuring a DNS server.

Changing the System Hostname

The hostname parameter is used to identify the system at the CLI prompt. You must enter a fully-qualified hostname for the system. The `sethostname` command sets the name of the IronPort appliance.

The sethostname Command

```
example.com> sethostname

example.com> hostname.com

example.com> commit
```

Configuring DNS Server(s)

You can configure the DNS settings for your IronPort appliance using the Network > DNS page or using the `dnsconfig` command. Before you configure DNS, consider the following:

- Whether to use the Internet's DNS servers or your own, and which specific server(s) to use.
- Which interface to use for DNS traffic.

  You must use the interface that faces the DNS server. If the appliance uses the M1 interface for data traffic, use M1. If the appliance uses only the P1 interface for data traffic, use P1. If the appliance uses both the P1 and P2 interfaces, then use P1 for an internal DNS server, and use P2 for a DNS server on the Internet.
- The number of seconds to wait before timing out a reverse DNS lookup.
- Clearing the DNS cache.

Specifying DNS Servers

IronPort 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 can specify alternate servers to use for specific domains. Since an alternate DNS server applies to a single domain, it must be authoritative (provide definitive DNS records) for that domain.

Split DNS

AsyncOS supports split DNS where internal servers are configured for specific domains and external or root DNS servers are configured for other domains. If you are using your own internal server, you can also specify exception domains and associated DNS servers.

Using the Internet Root Servers

The IronPort AsyncOS DNS resolver is designed to accommodate the large number of simultaneous DNS connections.

Multiple Entries and Priority

For each DNS server you enter, you can specify a numeric priority. AsyncOS will attempt to use the DNS server with the priority closest to 0. If that DNS server is not responding AsyncOS will attempt to use the server at the next priority. If you specify multiple entries for DNS servers with the same priority, the system randomizes the list of DNS servers at that priority every time it performs a query. The system then waits a short amount of time for the first query to expire or "time out" and then increments with a slightly longer amount of time for subsequent servers. The amount of time depends on the exact 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 timeout increments are 5, 10, 45.

For example, four DNS servers with two configured at priority 0, one at priority 1, and one at priority 2:

Table 19-5 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 priority 0 servers are down, the priority 1 server (1.2.3.6) is used, and 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.

DNS Alert

If an alert with the message "Failed to bootstrap the DNS cache" is generated when an appliance is rebooted, it means that the system was unable to contact its primary DNS servers. This 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

You can use the Clear DNS Cache button on Network > DNS page, or the `dnsflush` command to clear all information in the DNS cache when changes have been made to your local DNS system. Using this command might cause a temporary performance degradation while the cache is repopulated.

Configuring DNS

To edit DNS Settings:

1. Navigate to the Network > DNS page.
2. Click Edit Settings. The Edit DNS page appears.

Figure 19-17 Edit DNS Settings

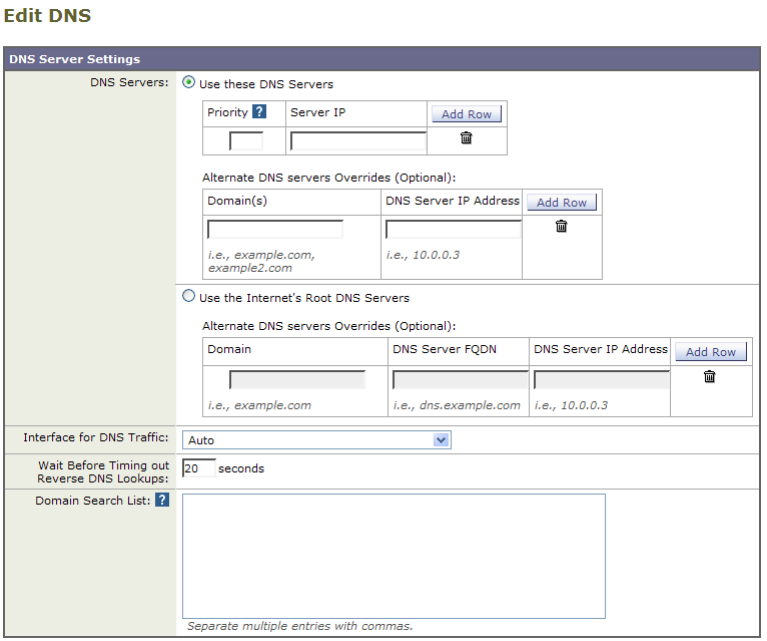

3. Select to use the Internet's root DNS servers or your own internal DNS server or the Internet's root DNS servers and specify authoritative DNS servers.
4. If you use your own DNS server(s), or specify authoritative DNS servers, enter the server ID, specify a priority, and use the Add Row key to repeat as necessary for each server.
5. Specify an interface for DNS traffic.

   Note — If the appliance uses both the P1 and P2 interfaces, you must use the interface that faces the DNS server. If you use an internal DNS server, use P1. If you use a DNS server on the Internet, use P2.

6. Enter the number of seconds to wait before cancelling a reverse DNS lookup.
7. Submit and commit to save the changes.

Configuring TCP/IP Traffic Routes

You can administer routes for data and management traffic, add static routes, load your IP routing tables, and modify the default gateway using the Network > Routes page or the `routeconfig` command.

The number of sections on this page depend on how you configured the "Restrict M1 port to appliance management services only" check box on the Network > Interfaces page:

- **Enabled.** When you use the Management interface for management traffic only, then this page includes two sections to enter gateway and static route table information, one for management traffic and one for data traffic. AsyncOS uses the management route information for management and data traffic, and data route information for data traffic. Figure 19-19 on page 397 shows the Routes page when the option is enabled.
- **Disabled.** When you use the Management interface for both management and data traffic, then this page includes one section to enter gateway and static route table information. AsyncOS uses the route information for both management and data traffic.

Note — A route gateway must reside on the same subnet as the Management or Data interface on which it is configured.

Modifying the Default Route

To modify the default gateway, use the Network > Routes page (or the `setgateway` CLI command), and click on Default Route in the corresponding table.

Figure 19-18 Editing the Default Route

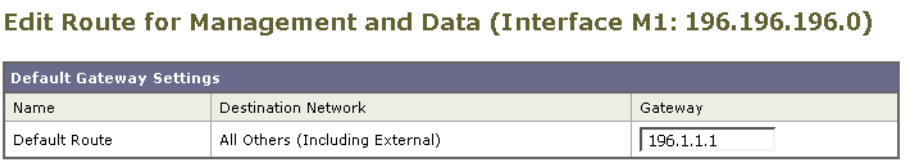

Enter the new gateway and click **Submit**. Commit your changes.

Working With Routing Tables

You can save your current routing table to a file. You can load a previously saved route table. You can add new routes or delete existing ones.

Saving and Loading Route Tables

To save a route table, click **Save Route Table** and specify where to save the file.

To load a previously saved route table, click **Load Route Table** and navigate to the file and then click **Submit**.

Adding a Route

To add a route:

1. Navigate to the Network > Routes page.

   Figure 19-19 Adding a Route

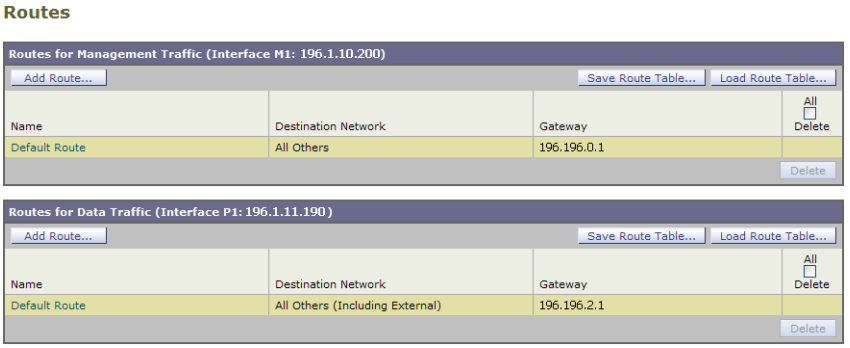

2. Click the **Add Route** button corresponding to the interface for which you are creating the route. The Add Route page is displayed.
3. Enter a Name, Destination Network, and Gateway.
4. Submit and commit your changes.

CONFIGURING NETWORK INTERFACES

You can configure the appliance network interfaces by modifying IP address, subnet, and host name information for the Management, Data, and L4 Traffic Monitor interfaces. Table 19-6 describes the network interface settings you can configure.

Table 19-6 Web Security Appliance Network Interface Settings

Interface	Port Number	Description
Management	M1	By default, the Management interface is used to administer the appliance and Web Proxy (data) monitoring. You can configure the M1 port for management only.
Data	P1 and P2 (proxy)	The Data interfaces are used for Web Proxy monitoring and L4 Traffic Monitor blocking (optional). You can also configure these interfaces to support outbound services such as DNS, software upgrades, NTP, and traceroute data traffic. For more information about configuring the Data interfaces, see "Configuring the Data Interfaces" on page 398.
L4 Traffic Monitor	T1 and T2	The L4 Traffic Monitor interfaces are used to configure a duplex or simplex wiring type. • **Duplex.** The T1 interface receives incoming and outgoing traffic. • **Simplex.** T1 receives outgoing traffic and T2 receives incoming traffic.

Note — If the Management and Data interfaces are all configured, each must be assigned IP addresses on different subnets.

You can manage the network interfaces using the following methods:

- **Web interface.** Use the Network > Interfaces page. For more information, see "Configuring the Network Interfaces from the Web Interface" on page 399.
- **Command line interface.** Use the `ifconfig` CLI command to create, edit, and delete network interfaces.

Configuring the Data Interfaces

You can configure the Web Security appliance to use any of the following combination of network interfaces for data traffic:

- **M1.** You can choose to use the Management interface for data traffic.
- **P1 only.** If you enable the P1 port, but not P2, you can use P1 for both incoming and outgoing data traffic.

- **P1 and P2.** If you enable both the P1 and P2 ports, you can use P1 for incoming data traffic and P2 for outgoing data traffic. When you enable both P1 and P2, you must use P2 for outgoing traffic and connect it to the Internet.

You can enable the M1 and P1 ports during or after System Setup. However, you can only enable the P2 port after System Setup using the `ifconfig` CLI command.

The Web Proxy listens for web requests on different network interfaces depending on how you configure the Web Security appliance:

- **M1.** The Web Proxy listens for requests on this interface when it is not configured to be restricted to appliance management services only.
- **P1.** The Web Proxy listens for requests on this interface when it is enabled.
- **P2.** By default, the Web Proxy does not listen for requests on this interface, even when enabled. However, you can configure it to listen for requests on P2 using the `advancedproxyconfig > miscellaneous` CLI command.

To configure the appliance to use P2 as a second data interface:

1. Configure the appliance to use P1 as the interface for data traffic. You can do this during System Setup or after initial setup on the Network > Interfaces page.
2. Use the `ifconfig` CLI command to enable P2. Use a different subnet for P2 than for P1.

   Note — If the Management and Data interfaces are all configured, each must be assigned IP addresses on different subnets.

3. In the web interface, go to the Network > Routes page. Change the data default rule to specify the next IP address that the P2 interface is connected to.
4. In the web interface, change the network interface for the following components to use P2 instead of P1: DNS and NTP servers.

   For more information, see "Configuring DNS Server(s)" on page 393 and "Editing System Time" on page 409.

Configuring the Network Interfaces from the Web Interface

To configure the network interfaces from the web interface:

1. Navigate to the Network > Interfaces page. Click **Edit Settings**.

   The Edit Interfaces page appears.

Figure 19-20 Editing Network Interfaces

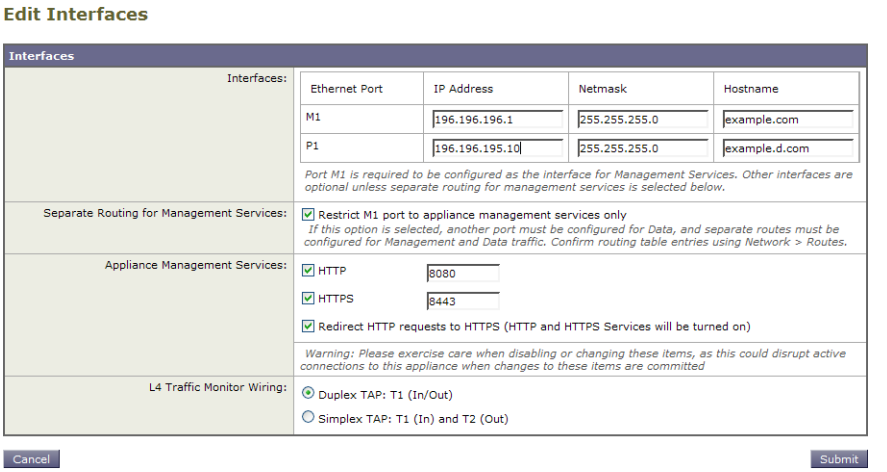

2. Configure interface settings as necessary.

   Table 19-7 describes the interface settings you can define for each interface.

Table 19-7 Interface Settings

Interface Setting	Description
IP Address	Enter the IP address to use to manage the Web Security appliance. Enter an IP address that exists on your management network.
Netmask	Enter the network mask to use when managing the Web Security appliance on this network interface.
Hostname	Enter the hostname to use when managing the Web Security appliance on this network interface.

Note — You can only enable and configure the P1 network interface for data traffic in the web interface. If you want to enable the P2 interface, you must use the `ifconfig` command. For more information about configuring the P2 interface, see "Configuring the Data Interfaces" on page 398.

3. Specify whether or not to have separate routing for the Management Services.

   If this checkbox is selected, the M1 port is restricted to appliance management services only. You will need to configure another port for data as well as separate routes for management and data traffic. For more information about configuring routes, see "Configuring TCP/IP Traffic Routes" on page 396.

4. Configure Appliance Management Services.

   Choose whether or not to use HTTP or HTTPS to administer AsyncOS through the web interface. You must specify the port to access AsyncOS with each protocol you configure.

   You can also choose to redirect HTTP requests to HTTPS. When you do this, AsyncOS automatically enables both HTTP and HTTPS.

5. Choose the type of wired connections plugged into the "T" network interfaces:

   - **Duplex TAP.** Choose Duplex TAP when the T1 port receives both incoming and outgoing traffic. You can use half- or full-duplex Ethernet connections.
   - **Simplex TAP.** Choose Simplex TAP when you connect the T1 port to the internal network (traffic flows from the clients to the Internet) and you connect the T2 port to the external network (traffic flows from the Internet to the clients).

   Note — IronPort recommends using simplex when possible because it can increase performance and security.

6. Submit and commit your changes.

CONFIGURING TRANSPARENT REDIRECTION

When you configure the Web Security appliance web proxy service in transparent mode, you must connect the appliance to an L4 switch or a WCCP v2 router, and you must configure the appliance so it knows to which device it is connected. You configure the device on the Network > Transparent Redirection page.

Figure 19-21 Network > Transparent Redirection Page

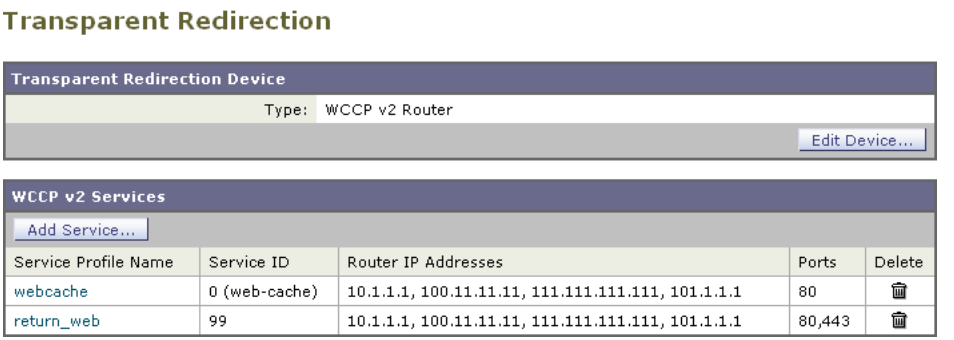

On this page, you can choose the device that transparently redirects traffic to the appliance, either an L4 switch or a WCCP router. When you choose an L4 switch as the device, there is nothing else to configure on this page.

However, when you choose a WCCP router as the device, you must create at least one WCCP service.

Working with WCCP Services

A WCCP service is an appliance configuration that defines a service group to a WCCP v2 router. It includes information such as the service ID and ports used. Service groups allow a web proxy to establish connectivity with a WCCP router and to handle redirected traffic from the router.

You can create WCCP services that use the following service types:

- **Standard service.** The standard service is also known as a well known service because the characteristics of it are known by both WCCP routers and the appliance. It redirects traffic on port 80. It is identified as the “web-cache” service.
- **Dynamic service.** Dynamic services are any other service a web proxy creates, but the web proxy must describe the components of the service group to the router. AsyncOS supports the creation of any dynamic service you choose to define. To create a dynamic service, you must provide the service ID number, port numbers, and specify whether to redirect packets based on the destination or source port and whether to distribute packets based on the client or server address.

The Web Cache Communication Protocol allows 257 different service IDs. AsyncOS allows you to create a dynamic WCCP service for each possible service ID. However, in typical usage, most users create one or two WCCP services, where one is a standard service and the other a dynamic service.

When you create a WCCP service of any type, you must also specify the following information:

- **Assignment method.** For more information, see "Working with the Assignment Method" on page 403.
- **Forwarding and Return method.** For more information, see "Working with the Forwarding and Return Method" on page 404.

If you enable IP spoofing on the appliance, you must create two WCCP services. For more information, see "IP Spoofing when Using WCCP" on page 404.

Working with the Assignment Method

WCCP defines the assignment method as the method by which redirected packets are distributed between web proxies. In this case, between one or more Web Security appliances. The assignment method determines how the router performs load balancing of packets among multiple Web Security appliances.

You configure the assignment method for a WCCP service in the Load-Balancing Method field under the Advanced section when you create or edit a WCCP service.

You can configure WCCP services to use either of the following assignment methods:

- **Mask.** This method relies on masking to make redirection decisions. WCCP routers make decisions using hardware in the router. This method can be very efficient because the hardware redirects the packets. You might want to choose mask to reduce CPU cycles on the router which can increase router performance. You can only use mask with WCCP routers that support mask assignment.

  Note — AsyncOS chooses the mask value to use with the router. You cannot configure the mask value.

- **Hash.** This method relies on a hash function to make redirection decisions. You might want to use Hash when the WCCP router does not support masking.

You can also configure a WCCP service to allow either mask or hash load balancing. When a WCCP service allows both mask and hash, AsyncOS communicates with the router to determine whether or not the router supports mask. If the router supports mask, then AsyncOS uses masking in the service group, if the router does not support mask, then AsyncOS uses hashing in the service group.

Working with the Forwarding and Return Method

WCCP defines the forwarding method as the method by which redirected packets are transported from the router to the web proxy. Conversely, the return method redirects packets from the web proxy to the router.

You configure the forwarding and return methods for a WCCP service in the Forwarding Method and Return Method fields under the Advanced section when you create or edit a WCCP service.

You can configure WCCP services to use either of the following methods:

- **Layer 2 (L2).** This method redirects traffic at layer 2 by replacing the packet's destination MAC address with the MAC address of the target web proxy. This method requires that the target web proxy be directly connected to the router at layer 2. WCCP routers only allow L2 negotiation when the appliance is directly connected to the router at layer 2. The L2 method redirects traffic at the router hardware level, and typically has better performance than Generic Routing Encapsulation (GRE). You might want to choose L2 when the router is directly connected to the appliance and you want the performance improvement provided by the L2 method. You can only use the L2 method with WCCP routers that support L2 forwarding.
- **Generic Routing Encapsulation (GRE).** This method redirects traffic at layer 3 by encapsulating the IP packet with a GRE header and a redirect header. This method redirects traffic at the router software level, which can impact performance. You might want to choose GRE when the appliance is not directly connected to the router.

You can also configure a WCCP service to allow either the L2 or GRE methods. When a WCCP service allows both L2 and GRE, the appliance uses the method that the router says it supports. If both the router and appliance support L2 and GRE, the appliance uses L2.

Note — If the router is not directly connected to the appliance, you must choose GRE.

IP Spoofing when Using WCCP

You can configure the Web Proxy to do IP spoofing. When enabled, requests originating from a client retain the client's source address and appear to originate from the client instead of the Web Proxy.

When you enable IP spoofing, you must create two WCCP services. One WCCP service must redirect traffic based on the destination port, and another based on the source port for the return path. The service based on the destination port can be the standard web-cache service. However, you must still create at least one dynamic service.

The two WCCP services you define for IP spoofing must have the same values for the following settings:

- Port numbers
- Router IP addresses
- Router security and password

Note — IronPort suggests using a service ID number from 90 to 99 for the WCCP service used for the return path (based on the source port).

For more information about creating WCCP services, see "Adding and Editing a WCCP Service" on page 405.

Adding and Editing a WCCP Service

You must create at least one WCCP service when you configure the transparent redirection device as a WCCP router. If IP spoofing is enabled on the appliance, you must create two WCCP services. For more information about IP spoofing, see "IP Spoofing when Using WCCP" on page 404.

To add or edit a WCCP service:

1. Navigate to the Network > Transparent Redirection page.

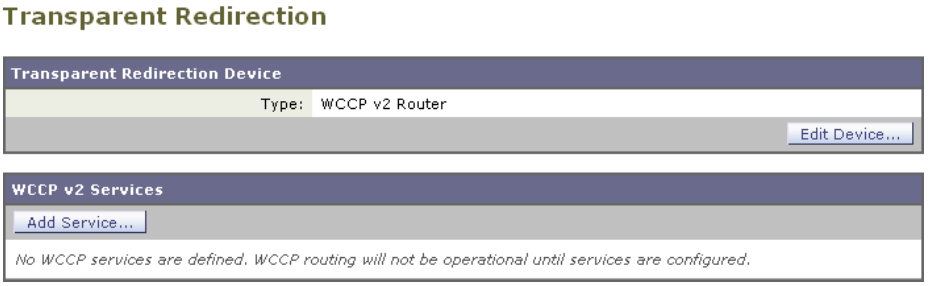

2. Verify the transparent redirection device is a WCCP v2 router. If it is not, click **Edit Device** to change it.
3. To add a WCCP service, click **Add Service**. Or, to edit a WCCP service, click the name of the WCCP service in the Service Profile Name column.

   The Add WCCP v2 Service page or Edit WCCP v2 Service page appears.

Add WCCP v2 Service

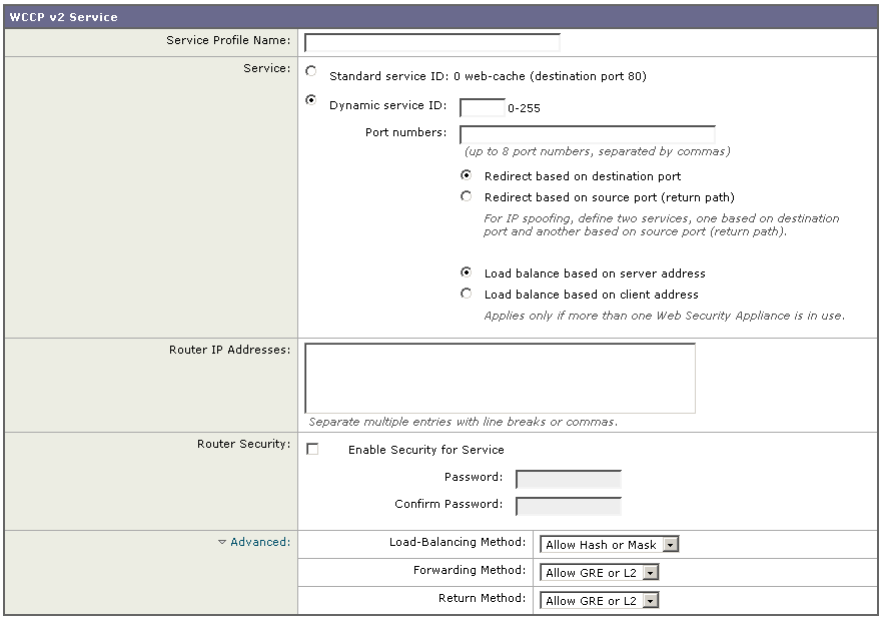

4. Configure the WCCP options.

   Table 19-8 describes the WCCP options.

Table 19-8 WCCP Service Options

WCCP Service Option	Description
Service Profile Name	Enter a name for the WCCP service.

Table 19-8 WCCP Service Options (Continued)

WCCP Service Option	Description
Service	Use this section to describe the service group for the router. Choose to create either a standard (“well known”) or dynamic service group. If you create a dynamic service, enter the following information: • **Service ID.** Enter any number from 0 to 255 in the Dynamic Service ID field. • **Port number(s).** Enter up to eight port numbers for traffic to redirect in the Port Numbers field. • **Redirection basis.** Choose to redirect traffic based on the source or destination port. Default is destination port. • **Load balancing basis.** When the network uses multiple Web Security appliances, you can choose how to distribute packets among the appliances. You can distribute packets based on the server or client address. When you choose client address, packets from a client always get distributed to the same appliance. Default is server address. For more information about well known and dynamic service groups, see “Working with WCCP Services” on page 402.
Router IP Addresses	Enter the IP address for one or more WCCP enabled routers. You can enter up to 32 routers to the service group. You must enter the IP address of each router. You cannot enter a multicast address.
Router Security	Choose whether or not to require a password for this service group. If required, enter the password in the password fields. The password can contain up to seven characters. When you enable security for a service group, every appliance and WCCP router that uses the service group must use the same password. Requiring a password enables you to control which routers and WCCP-enabled systems, such as the Web Security appliance, become part of the service group. WCCP uses the MD5 hash protocol to encrypt the password. **Note —** Each appliance or WCCP router in the service group authenticate the security component in a received WCCP packet immediately after validating the WCCP message header. Packets failing authentication are discarded.

Table 19-8 WCCP Service Options (Continued)

WCCP Service Option	Description
Advanced	Configure the following fields: • **Load-Balancing Method.** This is also known as the assignment method. Choose Mask, Hash, or both. Default is both. For more information about load-balancing, see "Working with the Assignment Method" on page 403. • **Forwarding Method.** Choose L2, GRE, or both. Default is both. For more information about the forwarding method, see "Working with the Forwarding and Return Method" on page 404. • **Return Method.** Choose L2, GRE, or both. Default is both. For more information about the return method, see "Working with the Forwarding and Return Method" on page 404.

5. Submit and commit your changes.

Deleting a WCCP Service

To delete a WCCP service:

1. Navigate to the Network > Transparent Redirection page.
2. Click the icon in the Delete column for the WCCP service you want to delete.
3. **Commit** your changes.

SETTING SYSTEM TIME

To set the system time on your Web Security appliance, set the time zone used, or select an NTP server and query interface. To set the system time, use the System Administration > Time Zone or Time Settings page or use the `ntpconfig`, `settime`, and `settz` commands.

Selecting a Time Zone

To set the time zone use the System Administration > Time Zone page:

Figure 19-22 The Time Zone Page

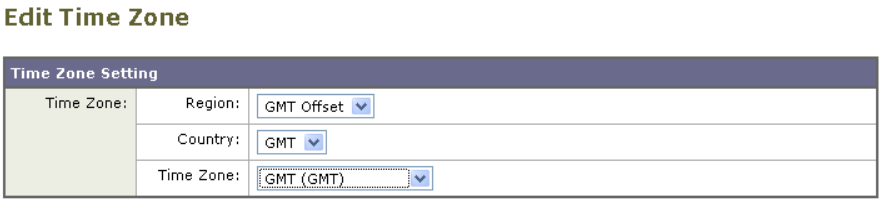

Select a time zone in the Time Zone area. You can configure the time zone by specifying the region and country, or by using a GMT offset.

The web interface uses the POSIX-style method of indicating the time zone using a GMT offset. This may be different than the offset convention used elsewhere.

The offset refers to the amount of hours that must be added or subtracted to the local time zone in order to reach GMT (Greenwich Mean Time or the Prime Meridian). Hours preceded by a minus sign ("-") are *east* of the Prime Meridian. A plus sign ("+") indicates *west* of the Prime Meridian.

For example, if the current time in New York is 08:00, then you must add five hours to get the current time in Greenwich, England, which is 13:00. In this case, to indicate the time in New York, the GMT offset is GMT+5. The "+5" in the offset indicates that you must add five hours to the time in New York to reach Greenwich Mean Time.

Editing System Time

To edit system time, use the System Administration > Time Settings page.

Figure 19-23 The Edit Time Settings Page

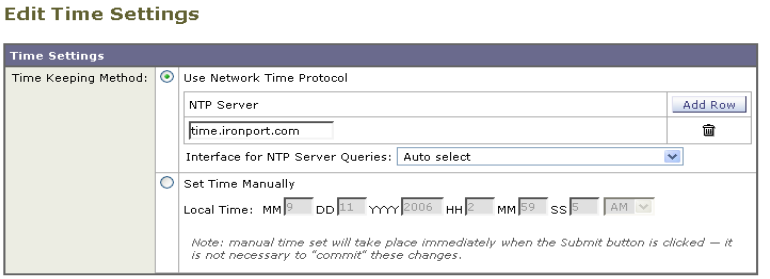

Configure NTP (Network Time Protocol)

To edit NTP server settings and use an NTP server to synchronize the system clock with other computers:

1. Enter an NTP server IP address and use the Add Row key to repeat as necessary for each NTP server.
2. Select an interface for NTP queries. This is the IP address from which NTP queries should originate.

   Note — If the appliance uses both the P1 and P2 interfaces, you must use the interface that faces the NTP server. If you use an internal NTP server, use P1. If you use a NTP server on the Internet, use P2.

3. Submit and commit the changes.

Manually Setting System Time

To set the system time manually:

1. Select Set Time Manually.
2. Enter the month, day, year, hour, minutes, and seconds.
3. Select A.M or P.M.
4. Submit and commit to save the changes.

INSTALLING A SERVER DIGITAL CERTIFICATE

When an administrator logs into the Web Security appliance using HTTPS, or when the appliance is configured for secure client authentication, the appliance uses a digital certificate to securely establish the connection with the client application. The Web Security appliance uses the "IronPort Appliance Demo Certificate" that comes installed by default. However, client applications are not programmed to recognize this certificate, so you can upload a digital certificate to the appliance that your applications recognize automatically.

For information about secure client authentication, see "Sending Authentication Credentials Securely" on page 279.

Figure 19-24 shows the warning message that is displayed in Firefox when accessing the Web Security appliance using the IronPort Appliance Demo Certificate.

Figure 19-24 IronPort Appliance Demo Certificate as an Unknown Authority

To configure the Web Security appliance to use a different digital server certificate, follow these steps:

1. Obtain a certificate and private key pair to upload. For more information, see "Obtaining Certificates" on page 411.
2. Upload the certificate and private key pair to the appliance. For more information, see "Uploading Certificates to the Web Security Appliance" on page 412.

Obtaining Certificates

To obtain a digital certificate to upload to the appliance, you must follow these steps:

1. Generate a public-private key pair.

2. Generate a Certificate Signing Requests (CSR).
3. Contact a certificate authority (CA) to sign the certificate.

The certificate you upload to the appliance must meet the following requirements:

- It must use the X.509 standard.
- It must include a matching private key in PEM format. DER format is not supported.
- The private key must be unencrypted.

The Web Security appliance cannot generate Certificate Signing Requests (CSR). Therefore, to have a certificate created for the appliance, you must issue the signing request from another system. Save the PEM-formatted key from this system because you will need to install it on the appliance later.

You can use any UNIX machine with a recent version of OpenSSL installed. Be sure to put the appliance host name in the CSR. Use the guidelines at the following location for information on generating a CSR using OpenSSL:

`http://www.modssl.org/docs/2.8/ssl_faq.html#ToC28`

Once the CSR has been generated, submit it to a certificate authority (CA). The CA will return the certificate in PEM format.

If you are acquiring a certificate for the first time, search the Internet for "certificate authority services SSL server certificates," and choose the service that best meets the needs of your organization. Follow the service's instructions for obtaining an SSL certificate.

Note — You can also generate and sign your own certificate. Tools for doing this are included with OpenSSL, free software from `http://www.openssl.org`.

Intermediate Certificates

In addition to root certificate authority (CA) certificate verification, AsyncOS supports the use of intermediate certificate verification. Intermediate certificates are certificates issued by a trusted root CA which are then used to create additional certificates. This creates a chained line of trust. For example, a certificate may be issued by example.com who, in turn, is granted the rights to issue certificates by a trusted root CA. The certificate issued by example.com must be validated against example.com's private key as well as the trusted root CA's private key.

Uploading Certificates to the Web Security Appliance

To upload a digital certificate to the Web Security appliance, use the `certconfig` command.

The following example shows a certificate being uploaded. You can also add intermediate certificates from this command.

```
example.com> certconfig

Currently using the demo certificate/key for HTTPS management access.
```

```
Choose the operation you want to perform:
- SETUP - Configure security certificate and key.
[]> setup

Management (HTTPS):
paste cert in PEM format (end with '.'):
-----BEGIN CERTIFICATE-----
MIICLDCCAdYCAQAwDQYJKoZIhvcNAQEEBQAwgaAxCzAJBgNVBAYTAlBUMRMwEQYD
VQQIEwpRdWVlbnNsYW5kMQ8wDQYDVQQHEwZMaXNib2ExFzAVBgNVBAoTDk5ldXJv
bmlvLCBMZGEuMRgwFgYDVQQLEw9EZXNlbnZvbHZpbWVudG8xGzAZBgNVBAMTEmJy
dXR1cy5uZXVyb25pby5wdDEbMBkGCSqGSIb3DQEJARYMc2FtcG9AaWtpLmZpMB4X
DTk2MDkwNTAzNDI0M1oXDTk2MTAwNTAzNDI0M1owgaAxCzAJBgNVBAYTAlBUMRMw
EQYDVQQIEwpRdWVlbnNsYW5kMQ8wDQYDVQQHEwZMaXNib2ExFzAVBgNVBAoTDk5l
dXJvbmlvLCBMZGEuMRgwFgYDVQQLEw9EZXNlbnZvbHZpbWVudG8xGzAZBgNVBAMT
EmJydXR1cy5uZXVyb25pby5wdDEbMBkGCSqGSIb3DQEJARYMc2FtcG9AaWtpLmZp
MFwwDQYJKoZIhvcNAQEBBQADSwAwSAJBAL7+aty3S1iBA/+yxjxv4q1MUTd1kjNw
L4lYKbpzzlmC5beaQXeQ2RmGMTXU+mDvuqItjVHOK3DvPK7lTcSGftUCAwEAATAN
BgkqhkiG9w0BAQQFAANBAFqPEKFjk6T6CKTHvaQeEAsX0/8YHPHqH/9AnhSjrwuX
9EBc0n6bVGhN7XaXd6sJ7dym9sbsWxb+pJdurnkxjx4=
-----END CERTIFICATE-----
.

paste key in PEM format (end with '.'):
-----BEGIN RSA PRIVATE KEY-----
MIIBPAIBAAJBAL7+aty3S1iBA/+yxjxv4q1MUTd1kjNwL4lYKbpzzlmC5beaQXeQ
2RmGMTXU+mDvuqItjVHOK3DvPK7lTcSGftUCAwEAAQJBALjkK+jc2+iihI98riEF
oudmkNziSRTYjnwjx8mCoAjPWviB3c742eO3FG4/soi1jD9A5alihEOXfUzloenr
8IECIQD3B5+0l+68BA/6d76iUNqAAV8djGTzvxnCxycnxPQydQIhAMXt4trUI3nc
a+U8YL2HPFA3gmhBsSICbq2OptOCnM7hAiEA6Xi3JIQECob8YwkRj29DU3/4WYD7
WLPgsQpwo1GuSpECICGsnWH5oaeD9t9jbFoSfhJvv0IZmxdcLpRcpslpeWBBAiEA
6/5B8J0GHdJq89FHwEG/H2eVVUYu5y/aD6sgcm+0Avg=
-----END RSA PRIVATE KEY-----
.

Do you want add an intermediate certificate? [N]> N

Currently using custom certificate/key for HTTPS management access.

Choose the operation you want to perform:
- SETUP - Configure security certificate and key.
[]>

example.com> commit

Please enter some comments describing your changes:
[]> Installed certificate and key for HTTPS management.
```

```
Changes committed: Fri Sep 26 17:59:53 2008 GMT
```

CHAPTER 20

Command Line Interface

This chapter contains the following information:

THE COMMAND LINE INTERFACE OVERVIEW

The IronPort AsyncOS Command Line Interface (CLI) is an interactive interface designed to allow you to configure and monitor the Web Security appliance. The commands are invoked by entering the command name with or without any arguments. If you enter a command without arguments, the command prompts you for the required information.

The Command Line Interface is accessible using SSH or Telnet on IP interfaces that have been configured with these services enabled, or using terminal emulation software on the serial port. By default, SSH and Telnet are configured on the Management port.

USING THE COMMAND LINE INTERFACE

This section describes the rules and conventions of the AsyncOS Command Line Interface.

Accessing the Command Line Interface

Access to the CLI varies depending on the management connection method chosen while setting up the appliance. The factory default username and password are listed next. Initially, only the admin user account has access to the CLI. You can add other users with differing levels of permission after you have accessed the CLI for the first time using the admin account. The System Setup Wizard prompts you to change the password for the admin account.

You can also reset the admin account password at any time using the `passwd` command.

You can connect using one of the following methods:

- **Ethernet.** Start an SSH or Telnet session with the IP address of the Web Security appliance. The factory default IP address is 192.168.42.42. SSH is configured to use port 22. Telnet is configured to use port 23.
- **Serial connection.** Start a terminal session with the communication port on your personal computer that the serial cable is connected to.

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

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

For example:

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

Working with the Command Prompt

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

```
example.com>
```

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

For example:

```
example.com> routeconfig

Choose a routing table:
- MANAGEMENT - Routes for Management Traffic
- DATA - Routes for Data Traffic
[]>
```

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

```
example.com> setgateway

Warning: setting an incorrect default gateway may cause the current
connection
to be interrupted when the changes are committed.
Enter new default gateway:
[172.xx.xx.xx]>
```

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

Command Syntax

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

```
example.com> logconfig
```

Select Lists

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

For example:

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

Yes/No Queries

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

For example:

```
Do you want to enable the proxy? [Y]> Y
```

Subcommands

Some commands give you the opportunity to use subcommand directives such as `NEW`, `EDIT`, and `DELETE`. The `EDIT` and `DELETE` functions provide a list of previously configured values.

For example:

```
example.com> interfaceconfig
```

```
Currently configured interfaces:
1. Management (172.xxx.xx.xx/xx: example.com)

Choose the operation you want to perform:
- NEW - Create a new interface.
- EDIT - Modify an interface.
- DELETE - Remove an interface.
[]>
```

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

Escaping Subcommands

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

Command History

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

Completing Commands

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

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

Configuration Changes

You can make configuration changes while web operations proceed normally.

Configuration changes do not take effect until you complete the following steps:

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

Changes to configuration that have not been committed are recorded, but do not go into effect until you run the `commit` command. However, not all commands require the `commit` command to be run.

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

GENERAL PURPOSE CLI COMMANDS

This section describes the some basic commands you might use in a typical CLI session, such as committing and clearing changes. For a full list of commands, see "Web Security Appliance CLI Commands" on page 422.

Committing Configuration Changes

The `commit` command allows you to change configuration settings while other operations proceed normally. Changes are not actually committed until you receive confirmation and a timestamp. Exiting the CLI session, system shutdown, reboot, failure, or issuing the `clear` command clears changes that have not yet been committed.

Entering comments after the commit command is optional.

```
example.com> commit


Please enter some comments describing your changes:

[]> Changed "psinet" IP Interface to a different IP address
Changes committed: Wed Jan 01 12:00:01 2007
```

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

Clearing Configuration Changes

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

```
example.com> clear

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

Changes cleared: Wed Jan 01 12:00:01 2007
example.com>
```

Exiting the Command Line Interface Session

The `exit` command logs you out of the CLI application. Configuration changes that have not been committed are cleared.

```
example.com> exit

Configuration changes entered but not committed. Exiting will lose
changes.
Type 'commit' at the command prompt to commit changes.
```

```
Are you sure you wish to exit?  [N]> y
```

Seeking Help on the Command Line Interface

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

```
example.com> help
```

WEB SECURITY APPLIANCE CLI COMMANDS

The Web Security Appliance CLI supports a set of proxy and UNIX commands to access, upgrade, and administer the system.

Table 20-1 lists the Web Security appliance Command Line Interface commands.

Table 20-1 Web Security appliance Administrative Commands

Command	Description
`advancedproxyconfig`	Configure more advanced Web Proxy configurations, such as authentication and DNS parameters. For more information about the `advancedproxyconfig` command, see "Advanced Proxy Configuration" on page 71.
`adminaccessconfig`	You can configure the Web Security appliance to have stricter access requirements for administrators logging into the appliance. For more information about the adminaccessconfig command, see "Configuring Administrator Settings" on page 377.
`alertconfig`	Specify alert recipients, and set parameters for sending system alerts.
`certconfig`	Configure security certificates and keys.
`clear`	Clears pending configuration changes since last commit.
`commit`	Commits pending changes to the system configuration.
`createcomputerobject`	Creates a computer object at the location you specify.
`dnsconfig`	Configure DNS server parameters.
`dnsflush`	Flush DNS entries on the appliance.
`etherconfig`	Configure Ethernet port connections.
`featurekey`	Submits valid keys to activate licensed features. For more information, see "Feature Keys Page" on page 369.
`featurekeyconfig`	Automatically check for and update feature keys. For more information, see "Feature Key Settings Page" on page 370.
`grep`	Searches named input files for lines containing a match to the give pattern.
`help`	Returns a list of commands.

Table 20-1 Web Security appliance Administrative Commands (Continued)

`ifconfig`
or
`interfaceconfig`
	Configure and manage network interfaces including M1, P1, and P2. Displays currently configured interfaces, and provides an operations menu to create, edit, or delete interfaces.

`smtprelay`	Configure SMTP relay hosts for internally generated email. An SMTP relay host is required to receive system generated email and alerts. For more information about configuring SMTP relay hosts, see "Configuring SMTP Relay Hosts" on page 386.
`last`	Lists user-specific user information that includes ttys and hosts, in reverse time order or lists the users that are logged in at a specified date and time.
`loadconfig`	Load a system configuration file.
`logconfig`	Configure access to log files.
`mailconfig`	Mail the current configuration file to the address specified.
`nslookup`	Queries Internet domain name servers for information about specified hosts and domains or to print a list of hosts in a domain.
`ntpconfig`	Configure NTP servers. Displays currently configured interfaces, and provides an operations menu to add, remove, or set the interface from whose IP address NTP queries should originate.
`packetcapture`	Intercepts and displays TCP/IP and other packets being transmitted or received over the network to which the appliance is attached. For more information, see "Packet Capture" on page 365.
`passwd`	Set the password.
`pathmtudiscovery`	Enables or disables Path MTU Discovery. You might want to disable Path MTU Discovery if you need to packet fragmentation.
`ping`	Sends an ICMP ECHO REQUEST to the specified host or gateway.
`proxyconfig` <`enable` | `disable`>	Enables or disables the Web Proxy.
`proxystat`	Display web proxy statistics.
`quit, q, exit`	Terminates an active process or session.

Table 20-1 Web Security appliance Administrative Commands (Continued)

`reboot`	Flushes the file system cache to disk, halts all running processes, and restarts the system.
`reportingconfig`	Configure a reporting system.
`resetconfig`	Restores the configuration to factory defaults.
`rollovernow`	Roll over a log file.
`routeconfig`	Configure destination IP addresses and gateways for traffic. Displays currently configured routes, and provides an operations menu to create, edit, or delete, or clear entries.
`saveconfig`	Saves a copy of the current configuration settings to a file. This file can be used to restore defaults, if necessary.
`setgateway`	Configure the default gateway for the machine.
`sethostname`	Set the hostname parameter.
`setntlmsecuritymode`	Changes the security setting for the NTLM authentication realm to either “ads” or “domain.” When the setting is “domain,” the appliance joins the Active Directory domain with a domain security trust account, and when the setting is “ads,” it joins the domain as a native Active Directory member. Default is ads.
`settime`	Set system time.
`settz`	Displays the current time zone, and provides an operations menu to set a local time zone.
`showconfig`	Display all configuration values.
`shutdown`	Terminates connections and shuts down the system.
`snmpconfig`	Configure the local host to listen for SNMP queries and allow SNMP requests.
`sshconfig`	Configure hostname and host key options for trusted servers.
`status`	Displays system status.
`supportrequest`	Send the support request email to IronPort customer care. This includes system information and a copy of the master configuration. The e-mail address is “support@ironport.com”.

Table 20-1 Web Security appliance Administrative Commands (Continued)

`tail`	Displays the end of a log file. Command accepts log file name or number as parameters.
	`example.com> tail system_logs`
	`example.com> tail 9`

`techsupport`	Provides a temporary connection to allow IronPort Customer Care/Applications Engineering to access the system and assist in troubleshooting.
`testauthconfig`	Tests the authentication settings for a given authentication realm against the authentication servers defined in the realm. For more information about testing authentication settings, see "Testing Authentication Settings" on page 268.
`telnet`	Communicates with another host using the TELNET protocol.
`traceroute`	Traces IP packets through gateways and along the path to a destination host.
`upgrade`	Install an AsyncOS software upgrade.
`upgradeconfig`	Configure upgrade server connections.
`userconfig`	Configure system administrators.
`version`	Displays general system information, installed versions of system software, and rule definitions.
`webcache`	Examine or modify the contents of the proxy cache, or configure domains and URLs that the appliance never caches. Allows an administrator to remove a particular URL from the proxy cache or specify which domains or URLs to never store in the proxy cache. For more information, see "Web Proxy Cache" on page 56.
`who`	Displays who is logged into the system.
`whoami`	Displays user information.

APPENDIX A

IronPort End User License Agreement

This appendix contains the following section:

CISCO IRONPORT SYSTEMS, LLC SOFTWARE LICENSE AGREEMENT

NOTICE TO ALL USERS: CAREFULLY READ THE FOLLOWING LEGAL AGREEMENT ("AGREEMENT") FOR THE LICENSE OF THE SOFTWARE (AS DEFINED BELOW). BY CLICKING THE ACCEPT BUTTON OR ENTERING "Y" WHEN PROMPTED, YOU (EITHER AN INDIVIDUAL OR A SINGLE ENTITY, COLLECTIVELY, THE "COMPANY") CONSENT TO BE BOUND BY AND BECOME A PARTY TO THE FOLLOWING AGREEMENT BETWEEN CISCO IRONPORT SYSTEMS, LLC, A DELAWARE CORPORATION ("IRONPORT") AND COMPANY (COLLECTIVELY, THE "PARTIES"). BY CLICKING THE ACCEPT BUTTON OR ENTERING "Y" WHEN PROMPTED, YOU REPRESENT THAT (A) YOU ARE DULY AUTHORIZED TO REPRESENT YOUR COMPANY AND (B) YOU ACCEPT THE TERMS AND CONDITIONS OF THIS AGREEMENT ON BEHALF OF YOUR COMPANY, AND AS SUCH, AN AGREEMENT IS THEN FORMED. IF YOU OR THE COMPANY YOU REPRESENT (COLLECTIVELY, "COMPANY") DO NOT AGREE TO THE TERMS AND CONDITIONS OF THIS AGREEMENT, CLICK THE CANCEL BUTTON OR ENTER "N" WHEN PROMPTED AND PROMPTLY (BUT NO LATER THAT THIRTY (30) DAYS OF THE DELIVERY DATE, AS DEFINED BELOW) NOTIFY IRONPORT, OR THE RESELLER FROM WHOM YOU RECEIVED THE SOFTWARE, FOR A FULL REFUND OF THE PRICE PAID FOR THE SOFTWARE.

1. DEFINITIONS

1.1 "Company Service" means the Company's email or internet services provided to End Users for the purposes of conducting Company's internal business and which are enabled via Company's products as described in the purchase agreement, evaluation agreement, beta or pre-release agreement, purchase order, sales quote or other similar agreement between the Company and IronPort or its reseller ("Agreement") and the applicable user interface and IronPort's standard system guide documentation that outlines the system architecture and its interfaces (collectively, the "License Documentation").

1.2 "End User" means the employee, contractor or other agent authorized by Company to access to the Internet or use email services via the Company Service.

1.3 "Service(s)" means (i) the provision of the Software functionality, including Updates and Upgrades, and (ii) the provision of support by IronPort or its reseller, as the case may be.

1.4 "Software" means: (i) IronPort's proprietary software licensed by IronPort to Company along with IronPort's hardware products; (ii) any software provided by IronPort's third-party licensors that is licensed to Company to be implemented for use with IronPort's hardware products; (iii) any other IronPort software module(s) licensed by IronPort to Company along with IronPort's hardware products; and (iv) any and all Updates and Upgrades thereto.

1.5 "Updates" means minor updates, error corrections and bug fixes that do not add significant new functions to the Software, and that are released by IronPort or its third party licensors. Updates are designated by an increase to the Software's release number to the right of the decimal point (e.g., Software 1.0 to Software 1.1). The term Updates specifically excludes Upgrades or new software versions marketed and licensed by IronPort or its third party licensors as a separate product.

1.6 "Upgrade(s)" means revisions to the Software, which add new enhancements to existing functionality, if and when it is released by IronPort or its third party licensors, in their sole discretion. Upgrades are designated by an increase in the Software's release number, located to the left of the decimal point (e.g., Software 1.x to Software 2.0). In no event shall Upgrades include any new versions of the Software marketed and licensed by IronPort or its third party licensors as a separate product.

2. LICENSE GRANTS AND CONSENT TO TERMS OF DATA COLLECTION

2.1 License of Software. By using the Software and the License Documentation, Company agrees to be bound by the terms of this Agreement, and so long as Company is in compliance with this Agreement, IronPort hereby grants to Company a non-exclusive, non-sublicensable, non-transferable, worldwide license during the Term to use the Software only on IronPort's hardware products, solely in connection with the provision of the Company Service to End Users. The duration and scope of this license(s) is further defined in the License Documentation. Except as expressly provided herein, no right, title or interest in any Software is granted to the Company by IronPort, IronPort's resellers or their respective licensors. This license and any Services are co-terminus.

2.2 Consent and License to Use Data. Subject to Section 8 hereof, and subject to the IronPort Privacy Statement at http://www.IronPort.com/privacy.html, as the same may be amended from time to time by IronPort with notice to Company, Company hereby consents and grants to IronPort a license to collect and use the data from the Company as described in the License Documentation, as the same may be updated from time to time by IronPort ("Data"). To the extent that reports or statistics are generated using the Data, they shall be disclosed only in the aggregate and no End User identifying information may be surmised from the Data, including without limitation, user names, phone numbers, unobfuscated file names, email addresses, physical addresses and file content. Notwithstanding the foregoing, Company may terminate IronPort's right to collect and use Data at any time upon prior written or electronic notification, provided that the Software or components of the Software may not be available to Company if such right is terminated.

3. CONFIDENTIALITY. Each Party agrees to hold in confidence all Confidential Information of the other Party to the same extent that it protects its own similar Confidential Information (and in no event using less than a reasonable degree of care) and to use such Confidential Information only as permitted under this Agreement. For purposes of this Agreement "Confidential Information" means information of a party marked "Confidential" or information reasonably considered by the disclosing Party to be of a proprietary or confidential nature; provided that the Data, the Software, information disclosed in design reviews and any pre-production releases of the Software provided by IronPort is expressly designated Confidential Information whether or not marked as such.

4. PROPRIETARY RIGHTS; OWNERSHIP. Title to and ownership of the Software and other materials and all associated Intellectual Property Rights (as defined below) related to the foregoing provided by IronPort or its reseller to Company will remain the exclusive property of IronPort and/or its superior licensors. Company and its employees and agents will not remove or alter any trademarks, or other proprietary notices, legends, symbols, or labels

appearing on or in copies of the Software or other materials delivered to Company by IronPort or its reseller. Company will not modify, transfer, resell for profit, distribute, copy, enhance, adapt, translate, decompile, reverse engineer, disassemble, or otherwise determine, or attempt to derive source code for any Software or any internal data files generated by the Software or to create any derivative works based on the Software or the License Documentation, and agrees not to permit or authorize anyone else to do so. Unless otherwise agreed in writing, any programs, inventions, concepts, documentation, specifications or other written or graphical materials and media created or developed by IronPort or its superior licensors during the course of its performance of this Agreement, or any related consulting or professional service agreements, including all copyrights, database rights, patents, trade secrets, trademark, moral rights, or other intellectual property rights ("Intellectual Property Right(s)") associated with the performance of such work shall belong exclusively to IronPort or its superior licensors and shall, in no way be considered a work made for hire for Company within the meaning of Title 17 of the United States Code (Copyright Act of 1976).

5. LIMITED WARRANTY AND WARRANTY DISCLAIMERS

5.1 Limited Warranty. IronPort warrants to Company that the Software, when properly installed and properly used, will substantially conform to the specifications in the License Documentation for a period of ninety (90) days from the delivery date or the period set forth in the License Documentation, whichever is longer ("Warranty Period"). FOR ANY BREACH OF THE WARRANTY CONTAINED IN THIS SECTION, COMPANY'S EXCLUSIVE REMEDY AND IRONPORT'S ENTIRE LIABILITY, WILL BE PROMPT CORRECTION OF ANY ERROR OR NONCONFORMITY, PROVIDED THAT THE NONCONFORMITY HAS BEEN REPORTED TO IRONPORT AND/OR ITS RESELLER BY COMPANY WITHIN THE WARRANTY PERIOD. THIS WARRANTY IS MADE SOLELY TO COMPANY AND IS NOT TRANSFERABLE TO ANY END USER OR OTHER THIRD PARTY. IronPort shall have no liability for breach of warranty under this Section or otherwise for breach of this Agreement if such breach arises directly or indirectly out of or in connection with the following: (i) any unauthorized, improper, incomplete or inadequate maintenance or calibration of the Software by Company or any third party; (ii) any third party hardware software, services or system(s); (iii) any unauthorized modification or alteration of the Software or Services; (iv) any unauthorized or improper use or operation of the Software or Company's failure to comply with any applicable environmental specification; or (v) a failure to install and/or use Updates, Upgrades, fixes or revisions provided by IronPort or its resellers from time to time.

5.2 WARRANTY DISCLAIMER. THE EXPRESS WARRANTIES SET FORTH IN SECTION 5.1 OF THIS AGREEMENT CONSTITUTE THE ONLY PERFORMANCE WARRANTIES WITH RESPECT TO THE SOFTWARE OR SERVICES. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IRONPORT LICENSES THE SOFTWARE AND SERVICES HEREUNDER ON AN "AS IS" BASIS. EXCEPT AS SPECIFICALLY SET FORTH HEREIN, IRONPORT AND ITS SUPERIOR LICENSORS MAKE NO REPRESENTATIONS OR WARRANTIES OF ANY KIND, WHETHER EXPRESS, IMPLIED, OR STATUTORY (EITHER IN FACT OR BY OPERATION OF LAW), AND EXPRESSLY DISCLAIM ALL OTHER WARRANTIES, INCLUDING WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NEITHER IRONPORT NOR ITS THIRD PARTY LICENSORS

WARRANT THAT THE SOFTWARE OR SERVICES (1) IS FREE FROM DEFECTS, ERRORS OR BUGS, (2) THAT OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED, OR (3) THAT ANY RESULTS OR INFORMATION THAT IS OR MAY BE DERIVED FROM THE USE OF THE SOFTWARE WILL BE ACCURATE, COMPLETE, RELIABLE AND/OR SECURE.

6. LIMITATION OF LIABILITY. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT WILL EITHER PARTY BE LIABLE TO THE OTHER FOR ANY LOSS OF PROFITS, COSTS OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, LOSS OF BUSINESS, LOSS OF USE OR DATA, INTERRUPTION OF BUSINESS, OR FOR INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES OF ANY KIND, EVEN IF SUCH PARTY RECEIVED ADVANCE NOTICE OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL THE LIABILITY OF EITHER PARTY ARISING UNDER ANY PROVISION OF THIS AGREEMENT, REGARDLESS OF WHETHER THE CLAIM FOR SUCH DAMAGES IS BASED IN CONTRACT, TORT, OR OTHER LEGAL THEORY, EXCEED THE TOTAL AMOUNT PAID FOR THE SOFTWARE OR SERVICES DURING THE TWELVE (12) MONTHS PRIOR TO THE EVENT GIVING RISE TO SUCH LIABILITY.

7. TERM AND TERMINATION. The term of this Agreement shall be as set forth in the License Documentation (the "Term"). If IronPort defaults in the performance of any material provision of this Agreement or the License Documentation, then Company may terminate this Agreement upon thirty (30) days written notice if the default is not cured during such thirty (30) day period. If Company defaults in the performance of any material provision of this Agreement or the License Documentation, IronPort may terminate this Agreement upon thirty (30) days written notice if the default is not cured during such thirty (30) day notice and without a refund. This Agreement may be terminated by one Party immediately at any time, without notice, upon (i) the institution by or against the other Party of insolvency, receivership or bankruptcy proceedings or any other proceedings for the settlement of such Party's debts, (ii) such other Party making a general assignment for the benefit of creditors, or (iii) such other Party's dissolution. The license granted in Section 2 will immediately terminate upon this Agreement's termination or expiration. Within thirty (30) calendar days after termination or expiration of this Agreement, Company will deliver to IronPort or its reseller or destroy all copies of the Software and any other materials or documentation provided to Company by IronPort or its reseller under this Agreement.

8. U.S. GOVERNMENT RESTRICTED RIGHTS; EXPORT CONTROL. The Software and accompanying License Documentation are deemed to be "commercial computer software" and "commercial computer software documentation," respectively, pursuant to DFAR Section 227.7202 and FAR Section 12.212, as applicable. Any use, modification, reproduction, release, performance, display or disclosure of the Software and accompanying License Documentation by the United States Government shall be governed solely by the terms of this Agreement and shall be prohibited except to the extent expressly permitted by the terms of this Agreement. Company acknowledges that the Software and License Documentation must be exported in accordance with U.S. Export Administration Regulations and diversion contrary to U.S. laws is prohibited. Company represents that neither the United States Bureau of Export Administration nor any other federal agency has suspended, revoked or denied Company export privileges. Company represents that Company will not use or transfer the

Software for end use relating to any nuclear, chemical or biological weapons, or missile technology unless authorized by the U.S. Government by regulation or specific license. Company acknowledges it is Company's ultimate responsibility to comply with any and all import and export restrictions, and other applicable laws, in the U.S. or elsewhere, and that IronPort or its reseller has no further responsibility after the initial sale to Company within the original country of sale.

9. MISCELLANEOUS. This Agreement is governed by the laws of the United States and the State of California, without reference to conflict of laws principles. The application of the United Nations Convention of Contracts for the International Sale of Goods is expressly excluded. Nothing contained herein shall be construed as creating any agency, partnership, or other form of joint enterprise between the parties. Neither party shall be liable hereunder by reason of any failure or delay in the performance of its obligations hereunder (except for the payment of money) on account of (i) any provision of any present or future law or regulation of the United States or any applicable law that applies to the subject hereof, and (ii) interruptions in the electrical supply, failure of the Internet, strikes, shortages, riots, insurrection, fires, flood, storm, explosions, acts of God, war, terrorism, governmental action, labor conditions, earthquakes, or any other cause which is beyond the reasonable control of such party. This Agreement and the License Documentation set forth all rights for the user of the Software and is the entire agreement between the parties and supersedes any other communications with respect to the Software and License Documentation. The terms and conditions of this Agreement will prevail, notwithstanding any variance with the License Documentation or any purchase order or other written instrument submitted by a party, whether formally rejected by the other party or not. This Agreement may not be modified except by a written addendum issued by a duly authorized representative of IronPort, except that IronPort may modify the IronPort Privacy Statement at any time, in its discretion, via notification to Company of such modification that will be posted at http://www.IronPort.com/privacy.html. No provision hereof shall be deemed waived unless such waiver shall be in writing and signed by IronPort or a duly authorized representative of IronPort. If any provision of this Agreement is held invalid, the remainder of this Agreement shall continue in full force and effect. The parties confirm that it is their wish that this Agreement has been written in the English language only.

10. IRONPORT CONTACT INFORMATION. If Company wants to contact IronPort for any reason, please write to IronPort Systems, Inc., 950 Elm Avenue, San Bruno, California 94066, or call or fax us at tel: 650.989.6500 and fax: 650.989.6543.

Index

T

W