


Sourcefire SSL Appliance Administration & Deployment Guide for SSL1500, SSL2000, and SSL8200

Software version: 3.6



Legal Notices

Cisco, the Cisco logo, Sourcefire, the Sourcefire logo, Snort, the Snort and Pig logo, and certain other trademarks and logos are trademarks or registered trademarks of Cisco and/or its affiliates in the United States and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.

The legal notices, disclaimers, terms of use, and other information contained herein (the "terms") apply only to the information discussed in this documentation (the "Documentation") and your use of it. These terms do not apply to or govern the use of websites controlled by Cisco or its subsidiaries (collectively, "Cisco") or any Sourcefire-provided or Cisco-provided products. Sourcefire and Cisco products are available for purchase and subject to a separate license agreement and/or terms of use containing very different terms and conditions.

The copyright in the Documentation is owned by Cisco and is protected by copyright and other intellectual property laws of the United States and other countries. You may use, print out, save on a retrieval system, and otherwise copy and distribute the Documentation solely for non-commercial use, provided that you (i) do not modify the Documentation in any way and (ii) always include Cisco's copyright, trademark, and other proprietary notices, as well as a link to, or print out of, the full contents of this page and its terms.

No part of the Documentation may be used in a compilation or otherwise incorporated into another work or with or into any other documentation or user manuals, or be used to create derivative works, without the express prior written permission of Cisco. Cisco reserves the right to change the terms at any time, and your continued use of the Documentation shall be deemed an acceptance of those terms.

© 2004 - 2014 Cisco and/or its affiliates. All rights reserved.

Disclaimers

THE DOCUMENTATION AND ANY INFORMATION AVAILABLE FROM IT MAY INCLUDE INACCURACIES OR TYPOGRAPHICAL ERRORS. CISCO MAY CHANGE THE DOCUMENTATION FROM TIME TO TIME. CISCO MAKES NO REPRESENTATIONS OR WARRANTIES ABOUT THE ACCURACY OR SUITABILITY OF ANY CISCO-CONTROLLED WEBSITE, THE DOCUMENTATION AND/OR ANY PRODUCT INFORMATION. CISCO-CONTROLLED WEBSITES, THE DOCUMENTATION AND ALL PRODUCT INFORMATION ARE PROVIDED "AS IS" AND CISCO DISCLAIMS ANY AND ALL EXPRESS AND IMPLIED WARRANTIES, INCLUDING BUT NOT LIMITED TO WARRANTIES OF TITLE AND THE IMPLIED WARRANTIES OF MERCHANTABILITY AND/OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL CISCO BE LIABLE TO YOU FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, PUNITIVE, OR CONSEQUENTIAL DAMAGES (INCLUDING BUT NOT LIMITED TO PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, LOSS OF DATA, LOSS OF PROFITS, AND/OR BUSINESS INTERRUPTIONS), ARISING OUT OF OR IN ANY WAY RELATED TO CISCO-CONTROLLED WEBSITES OR THE DOCUMENTATION, NO MATTER HOW CAUSED AND/OR WHETHER BASED ON CONTRACT, STRICT LIABILITY, NEGLIGENCE OR OTHER TORTUOUS ACTIVITY, OR ANY OTHER THEORY OF LIABILITY, EVEN IF CISCO IS ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES/JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF LIABILITY FOR CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATIONS MAY NOT APPLY TO YOU.

Table of Contents

Chapter 1:	Introduction to the Sourcefire SSL Appliances	8
	SSL Inspection Overview	9
	Product Overview	10
	Key Features	12
	Product Specifications	13
	Product Checklist	17
	Documentation Conventions	17
 Chapter 2:	 System Behavior & Deployment Examples.....	 19
	Transparent SSL Decryption / Encryption	19
	SSL Decryption Methods	20
	Known Server Key Method.....	21
	Certificate Re-Signing Method.....	23
	Self-Signed Server Certificate Handling.....	24
	Decryption Methods in Cooperative Configurations.....	25
	Marking SSL Plaintext.....	26
	Deployment Modes	27
	Passive-Tap Mode	29
	Passive-Inline Mode.....	32
	Active-Inline Mode	33

Table of Contents

	Policies.....	35
	Segment Policies	36
	Ruleset Policies.....	37
	Lists	46
	Reset Generation	46
	Failure Modes and High Availability	47
	Link Failures	48
	Software (Data-Plane) Failures.....	49
	Example Deployment Configurations	50
	Outbound Inspection	50
	Inbound Inspection	52
	Inbound and Outbound Inspection	52
	High Availability Deployment	54
Chapter 3:	Physical Installation	56
	Safety Information	56
	Requirements Checklist.....	57
	The following will be required:	
	Rack Mounting.....	57
	Rear	58
	Front Panel.....	59
	Connecting to the Network	64
Chapter 4:	Initial Configuration and Setup	67
	Bootstrap Phase	67
	Configuring Static IP Address for Management	69
	Password Entry.....	69
	Installation Process.....	73
	Network Connections	76
	Post Bootstrap Configuration.....	76
	Configuring System Date/Time and Time Zone	78
	Configuring Management Network Settings	79
	Configuring Management Users.....	80
	System Status.....	82
	Installing a CA for Certificate Re-Sign	83
	Creating a CA.....	83
	Importing a CA.....	86

Table of Contents

Importing Known Server Keys	86
Example Passive-Tap Mode Inspection	88
Example Passive-Inline Mode Inspection	94
Example Active-Inline Mode Inspection	98
Chapter 5:	
Web-Based Management Interface (WebUI)	102
Browser Configuration	102
Login Process	103
Screen Layout Explained	104
Monitoring the System	107
Dashboard	107
System Log	109
SSL Session Log	109
SSL Statistics	111
Invalid Certificates	112
Errors	113
Diagnostics	113
Debug	114
Configuring Segments and Policies	115
Rulesets	116
Segments	119
Distinguished Names List	123
Common Names List	124
IP Address Lists	125
Cipher Suites List	125
PKI Management	126
Internal Certificate Authorities	126
External Certificate Authorities	127
Certificate Revocation Lists	128
Trusted Certificates	129
Known Certificates and Keys	130
Known Keys List	131
Platform Management	131
Information	132
Management Network	133
Remote Logging	133
Date/Time	134

Table of Contents

Users	134
TACACS Servers	134
Alerts	135
Backup/Restore.....	136
Halt/Reboot.....	137
Import UI Certificate/Key	137
Update	138
Preferences.....	139
User Management.....	140
Change Password.....	140
Logout.....	140
Chapter 6:	Troubleshooting the System..... 141
Supported Network Protocols and Frame Encapsulations.....	141
Supported SSL/TLS Versions	142
Support for Client Certificates	142
Supported Cipher Suites.....	142
Support for SSL Record Layer Compression	145
Support for Stateless Session Resumption (RFC5077)	145
Steps to Troubleshoot SSL Decryption.....	145
Monitor Network Port Statistics	145
Monitor the SSL Statistics	146
Monitor the SSL Session Log	146
Verify Inspection Policy Configuration.....	146
Known Server Versus Trusted Server Certificates	146
Caveats when Enabling/Disabling SSL Inspection	146
Generating the Internal CA Certificates	147
Access to Microsoft Windows Update Denied.....	147
Issues with Alerts	148
Procedure for Reporting an Issue	148
Preparing for Hardware Diagnostics or Maintenance	148
Chapter 7:	Safety Information 149
Safety Instructions	149
Rack Mounting the Equipment	150
Chapter 8:	Licenses and Licensing Terms..... 151
GNU General Public License.....	151
OpenSSL License	156

Table of Contents

Copyright Statement.....	157
Limited Warranty	157
Liability.....	158
Disclaimer of Damages.....	158
Export Regulations.....	159
General	159
Excluded Software.....	160
Chapter 9:	
Technical Support.....	161

Chapter 1

Introduction to the Sourcefire SSL Appliances

As organizations become dependent on IP-based applications and services, the demand for secure, reliable communications has never been higher. The increase in CPU performance has made client-based encryption a viable solution for enterprise communications. The Secure Sockets Layer (SSL) cryptographic protocol is the dominant client-based encryption protocol and now constitutes a significant and growing percentage of the traffic in the enterprise LAN and WAN as well as throughout service provider networks. SSL is used as a Virtual Private Network (VPN) technology to allow users to securely communicate with the enterprise. It is also used for secure communications from inside of the enterprise to Internet-based applications and services (banking, e-commerce, web mail, cloud applications and personal email).

The privacy benefits provided by SSL can quickly be overshadowed by the risks it brings to the enterprise network. SSL encryption can:

- mask threats, such as viruses, spam and malware
- make corporate acceptable use policies less effective
- increase the likelihood of accidental or intentional leakage of confidential information

The Sourcefire SSL appliance decrypts SSL traffic up to 2Gbps to enable existing security appliances to effectively inspect SSL traffic. The SSL appliance operates transparently on the network and supports both passive and inline network configurations.

SSL Inspection Overview

The ability to inspect SSL traffic enables existing security and network appliances to access the plaintext within SSL flows, thereby enabling the security appliance to do its job. SSL inspection is a complex and computationally intensive process that can easily become a performance bottleneck unless implemented with appropriate hardware acceleration techniques.

There are two different mechanisms that can be used to inspect SSL traffic depending on what information is available and how the inspection device is deployed in the network.

- Known server key mechanism relies on the inspecting device having a copy of the servers private key and certificate
- Certificate re-sign mechanism relies on the inspecting device having a trusted CA certificate that can be used to sign SSL server certificates that have been intercepted and modified

There are three basic connectivity modes that define how the SSL inspecting appliance and the associated security appliance are connected to each other and to the network. These modes are identified as:

- Active-inline
- Passive-inline
- Passive-tap

The active/passive designation refers to the associated security appliance and how it behaves while the inline/tap designation refers to how the SSL inspecting device is connected to the network. An *active* associated appliance processes traffic from the SSL inspecting device and then returns the traffic to the device while a *passive* appliance simply consumes traffic. The SSL inspecting device can be either inline or can be connected to a network span or tap port.

IMPORTANT! SSL inspection using **certificate re-sign** and SSL policy enforcement can only be done if the SSL appliance is connected inline in the network.

IMPORTANT! Only **known server key** mode can be used to inspect SSL traffic when the inspecting device is connected to a network tap. Inspection will not be possible if the session uses Diffie-Hellman for key exchange.

SSL inspection enables the identification and elimination of risks, such as regulatory compliance violations, viruses or malware, and intrusion attempts normally hidden within SSL. The privacy and integrity of SSL-encrypted communications are maintained by making the plaintext available only to the directly attached appliance. This requires the environment to be physically secure.

Additional privacy for SSL-encrypted traffic can be achieved by configuring appropriate policies to control which traffic is inspected and which is not.

IMPORTANT! The SSL appliance and the associated security appliances that it is enabling to inspect traffic should all be located in a physically secure environment to prevent unauthorized access to the decrypted SSL traffic.

Product Overview

The Sourcefire SSL Appliance 1500, SSL Appliance 2000, and SSL Appliance 8200 are high performance transparent proxies for Secure Socket Layer (SSL) network communications. It enables a variety of applications to access the plaintext (that is, original unencrypted data) in SSL-encrypted connections and has been designed for security and network appliance manufacturers, enterprise IT organizations and system integrators. Without compromising any aspect of enterprise policies or government compliance, the SSL appliance allows network appliances to be deployed with highly granular flow analysis while maintaining line rate performance.

Sourcefire's SSL appliance product provides two main functions when deployed within a network:

- It enables other security appliances to see a non-encrypted version of SSL traffic that is crossing the network. This is called SSL inspection because the security appliance is able to inspect the decrypted traffic for possible threats, which is something it cannot do when it sees encrypted traffic.
- It can act as a *policy control point* enabling explicit control over what SSL traffic is and is not allowed across the network.

The SSL appliance is designed to work alongside existing security devices such as Intrusion Prevention Systems (IPS), Intrusion Detection Systems (IDS), Data Loss Prevention systems (DLP), Network Forensic appliances and so on. It provides a non-encrypted version of SSL traffic to the associated appliance while maintaining an end-to-end SSL connection between the client and server involved in the session.

Unlike most other SSL proxy devices, the SSL appliance does not rely on the IP destination port number being used by a session to determine if it is using SSL or not. The SSL appliance uses deep packet inspection to identify SSL flows, this ensures that it is capable of finding and inspecting any SSL traffic in the network even if the traffic is using non-standard port numbers.

The SSL appliance incorporates flow processing hardware and cryptographic acceleration hardware enabling it to forward non-SSL traffic at multiple gigabit per second rates, while offering industry-leading transparent proxy performance (that is, decrypting and re-encrypting) for SSL traffic.

The SSL appliance supports two different mechanisms that allow inspection of SSL. Each mechanism requires that different information is available to the SSL appliance.

- *Known server key mechanism* relies on the inspecting device having a copy of the SSL server's private key and certificate
- *Certificate re-sign mechanism* relies on the inspecting device having a trusted CA certificate that can be used to sign SSL server certificates that have been intercepted and modified

The mechanism used to inspect an SSL flow can be chosen based on the details related to that flow so it is possible for an SSL appliance to be configured to use both mechanisms at the same time.

The mechanism used to inspect an SSL flow can be chosen based on the details related to that flow so it is possible for an SSL Inspector Appliance to be configured to use both mechanisms at the same time.

There are three basic connectivity modes that define how the SSL Inspector Appliance and the associated security appliance are connected to each other and to the network. These modes are identified as:

- Active-Inline
- Passive-Inline
- Passive-Tap

The Active / Passive designation refers to the associated security appliance and how it behaves while the Inline/Tap designation refers to how the SSL appliance is connected to the network. An "Active" associated appliance processes traffic from the SSL appliance and then returns the traffic to the SSL appliance while a "Passive" appliance simply consumes traffic. The SSL appliance can be either "Inline" or can be connected to a network span or tap port.

IMPORTANT! SSL inspection using certificate re-sign and SSL policy enforcement can only be done if the SSL appliance is connected inline in the network.

It is possible to have more than one associated security appliance connected to an SSL appliance and receiving the inspected traffic. A typical configuration would be an IPS device attached to an SSL appliance operating in active-inline mode with a network forensic appliance also connected in passive mode and receiving the same data that is going through the IPS. The ability to mirror the output of the SSL appliance to additional passive appliances is a useful feature that removes the need for an external device to mirror traffic to more than one appliance.

The SSL appliance enables the identification and elimination of risks, such as regulatory compliance violations, viruses/malware, and intrusion attempts normally hidden within SSL. The privacy and integrity of SSL-encrypted communications are maintained by making the plaintext available only to the attached appliance. This requires the environment to be physically secure.

Additional privacy for SSL-encrypted traffic can be achieved by configuring appropriate policies to control which traffic is inspected.

The SSL appliance and the associated security appliances that it is enabling to inspect traffic should all be located in a physically secure environment to prevent unauthorized access to the decrypted SSL traffic.

The act of inspecting SSL traffic may be subject to corporate policy guidelines or national legislation. It is your responsibility to ensure that your use of the SSL appliance is in accordance with any such legal or policy requirements.

Key Features

The SSL Appliance 1500 (or SSL1500), SSL Appliance 2000 (or SSL2000), and SSL Appliance 8200 (or SSL8200) provide a complete solution to the problem of dealing with threats contained within encrypted SSL traffic. A single SSL appliance can be deployed to detect and inspect all SSL traffic that may pose a threat and can pass the decrypted content to one or more network security appliances which can record or block any threats. The ability to feed inspected traffic to more than one associated security appliance ensures that SSL traffic only has to be decrypted and then re-encrypted once as it crosses the network.

Line Rate Network Performance for 10, 100, GigE and 10G Links

All non-SSL traffic flows are cut through (that is, forwarded directly from port to port) by the embedded flow processor (NFP-3240), minimizing latency for traffic such as Voice over Internet Protocol (VoIP).

Network Transparency

The SSL appliance is deployed as a *bump-in-the-wire* and is completely transparent to both end systems and intermediate networking elements. There is no need for network reconfiguration, IP addressing or topology changes, or modifications to client or server software (for example, changing web proxy settings or client IP addresses).

Compatible with Existing Devices and Applications

Intercepted plaintext is delivered to attached devices as a valid regenerated TCP stream via the SSL appliance's network ports. This allows existing security appliances (such as IDS, IPS, firewall, lawful intercept, and compliance monitoring devices) to expand their scope to also provide benefits for SSL-encrypted traffic.

Supports Multiple Decryption Methods and Various Encryption Algorithms and Protocols

One decryption method supports situations where server keys can be obtained, while another method can decrypt traffic to servers on the Internet; therefore the SSL appliance supports both *inbound* as well as *outbound* SSL traffic. The SSL appliance can accommodate most SSL-encrypted protocols (for example, web (HTTPS), email protocols, and most other standard or proprietary protocols). Either SSL 3.0, TLS 1.0, TLS 1.1, or TLS 1.2 can be used.

High Availability Deployment Options

Link state mirroring and fail-to-wire/fiber options allow the SSL appliance to be deployed in configurations that ensure connectivity is maintained even if hardware fails or software is temporarily not fully functional (for example, because software is being upgraded).

Traffic Mirroring

The ability to mirror copies of the traffic on an interface to up to two other interfaces enables multiple network security appliances to receive the inspected traffic flows. For example, an IPS may be attached to the SSL appliance and at the same time a network forensics appliance could be connected with both appliances receiving the inspected traffic flows.

Traffic Aggregation

When the SSL appliance is used in tap mode (connected to a network tap rather than inline), it can be configured to aggregate traffic received on multiple interfaces onto a single logical segment which contains the policies for how the traffic should be processed. This avoids the need to use an external aggregation device when traffic is being collected from multiple network taps.

Product Specifications

Sourcefire offers three SSL appliances: the SSL Appliance 1500, SSL Appliance 2000, and SSL Appliance 8200. These three devices offer the same functionality and use the same software. The Sourcefire SSL Appliance 1500 is a 1U device with eight interfaces. The Sourcefire SSL Appliance 2000 is also a 1U device, but has three NetMod slots. The Sourcefire SSL Appliance 8200 is a larger unit which takes up 2U of rack space and can have as many as seven NetMods.

The specifications in the table may change over time. Any changes will be reflected in new versions of this documentation, which you can download from the Sourcefire support site.

SSL Appliance 1500 Specification

Category	Description
Chassis Dimensions	17.5" (W) x 19.5" (D) x 1.75" (H) (444.5mm x 495.3mm x 44.5mm)
Weight	29 lbs (13.15 kg)
Processors	2 x Intel Xeon E5620 quad core CPUs
System memory	16GB DDR3
Network Flow Engine (NFE)	1 x NFE-3240 card (NFP-3240 + 4GB DDR3 + PCIe gen2 x8)
Interfaces	8 x 10/100/1000 Ethernet interfaces
Management Network interfaces	2 x 10/100/1000 copper interfaces on rear panel
Integrated Display	16 character by 2 line LCD on front panel
Power Supplies	2 x 450W redundant hot swap power supplies
Operating Temperature	0°C to 40°C
Storage Temperature	-10°C to 70°C

The specifications in the table may change over time. Any changes will be reflected in new versions of this documentation, which you can download from the Sourcefire support site.

SSL Appliance 2000 Specification

Category	Description
Chassis Dimensions	17.2" (W) x 19.0" (D) x 3.48" (H) (433mm x 735mm x 88.2mm)
Weight	58 lbs (26.4 kg)
Processors	2 x Intel Xeon E5645 hex core CPUs
System memory	48GB DDR3
Network Flow Engine (NFE)	2 x NFE-3240 card (NFP-3240 + 4GB DDR3 + PCIe gen2 x8)
Network Module slots (NetMods)	7 x NetMod slots (system limit is total of 16 interfaces)
Supported NetMod types (all NetMods have fail-to-wire/open capabilities)	<ul style="list-style-type: none">• 2 x 10G fiber• 4 x 10/100/1000 fiber• 4 x 10/100/1000 copper
Management Network interfaces	2 x 10/100/1000 copper interfaces on rear panel
Integrated Display	16 character by 2 line LCD on front panel
Power Supplies	2 x 750W redundant hot swap power supplies
Operating Temperature	0°C to 40°C
Storage Temperature	-10°C to 70°C
Cooling	Generates up to 1725 BTU/hour worst case
Air flow	160ft ³ /min (4.5m ³ /min)

The specifications in the [SSL Appliance 8200 Specification](#) table may change over time. Any changes will be reflected in new versions of this documentation, which you can download from the Sourcefire support site.

SSL Appliance 8200 Specification

Category	Description
Chassis Dimensions	17.5" (W) x 19.5" (D) x 1.75" (H) (444.5mm x 495.3mm x 44.5mm)
Weight	29 lbs (13.15 kg)
Processors	2 x Intel Xeon E5620 quad core CPUs
System memory	24GB DDR3
Network Flow Engine (NFE)	1 x NFE-3240 card (NFP-3240 + 4GB DDR3 + PCIe gen2 x8)
Network Module slots (NetMods)	3 x NetMod slots – all NetMods must be same speed
Supported NetMod types (all NetMods have fail-to-wire/open capabilities)	<ul style="list-style-type: none">• 2 x 10G fiber• 4 x 10/100/1000 fiber• 4 x 10/100/1000 copper
Management Network interfaces	2 x 10/100/1000 copper interfaces on rear panel
Integrated Display	16 character by 2 line LCD on front panel
Power Supplies	2 x 650W redundant hot swap power supplies
Operating Temperature	0°C to 40°C
Storage Temperature	-10°C to 70°C
Cooling	Generates up to 2225 BTU/hour worst case
Air flow	210ft ³ /min (6m ³ /min)

Product Checklist

Carefully unpack the Sourcefire SSL appliance and compare the actual contents with the [Product Checklist](#) to ensure that you have received all ordered components. Follow the instructions in [Physical Installation](#) on page 56 and [Initial Configuration and Setup](#) on page 67 to install and initially configure the appliance.

Product Checklist

Part	Description	Quantity
Sourcefire SSL1500 or Sourcefire SSL2000 or Sourcefire SSL8200	1U rack-mountable device 2U rack-mountable device	1
Power Cords	One for each redundant power supply	2
Rack-mounting rails	Rails to rack mount the device	2
Administration and Deployment Guide	On documentation CD	1
Getting Started Guide	Paper copy and on documentation CD	1
Release Notes	Paper copy	1
Safety Notice	Single sheet safety notice	1

Documentation Conventions

The following conventions are used throughout this document.

IMPORTANT! A note provides additional information that may be of interest.

WARNING! A warning provides additional information that you need to pay attention to.

Throughout this document the term SSL is used to mean both SSL and TLS, unless explicitly indicated. Secure Socket Layer (SSL) has been largely replaced by Transport Layer Security (TLS) which is the more up-to-date standard derived

from SSL. Both SSL and TLS traffic are present in networks today and the Sourcefire SSL appliance is capable of inspecting both types of traffic.

IMPORTANT! The embedded software contained within the Sourcefire SSL Appliance 1500, SSL Appliance 2000, and SSL Appliance 8200 is subject to licensing terms and conditions imposed by Sourcefire and third party software providers. You should only use the SSL1500, SSL2000, or SSL8200 if you agree to these licensing conditions. See [Licenses and Licensing Terms](#) on page 151 for details of licensing terms and conditions.

IMPORTANT! The act of inspecting SSL traffic may be subject to corporate policy guidelines or national legislation. It is your responsibility to ensure that your use of the SSL appliance is in accordance with any such legal or policy requirements.

Chapter 2

System Behavior & Deployment Examples

This section describes the functions performed by the Sourcefire SSL Appliance 1500, SSL Appliance 2000, and SSL Appliance 8200, its behavior, and its interaction with attached devices. For details on how to setup and configure the SSL appliance see [Initial Configuration and Setup](#) on page 67 and [Web-Based Management Interface \(WebUI\)](#) on page 102.

Transparent SSL Decryption / Encryption

The main function of the Sourcefire SSL appliance is to decrypt SSL traffic to obtain the plaintext sent within the SSL-encrypted session. The plaintext information is fed to one or more attached devices for processing or analysis. Because the plaintext data stream is repackaged as a valid TCP stream, applications that are hosted on the attached devices do not need to be modified to process the received plaintext stream.

The Sourcefire SSL appliance provides SSL inspection capabilities to existing devices.

A network *segment* is the collection of SSL appliance interfaces used to connect to the network carrying the traffic that is being inspected and to the attached appliances that are processing the traffic. Depending on how the SSL appliance is connected and on how many attached appliances are connected a segment may contain up to eight interfaces.

When used in active-inline mode or passive-inline mode the SSL appliance acts as a fully transparent proxy: the Ethernet ports used to connect it to the data network do not have IP addresses, and the other devices in the network are

unaware that the SSL appliance has been installed. No changes to clients or other network equipment are required when installing the SSL appliance.

If used in active-inline mode or passive-inline mode, the SSL appliance is a layer 2 *bump-in-the-wire* device and can be deployed without renumbering the existing IP network. In most cases, no network topology changes are required.

If used in passive-tap mode, the SSL appliance is no longer a bump-in-the-wire on the live network, but rather a bump-in-the-wire on the passive link between the network SPAN/tap device and the attached appliances.

The SSL appliance can detect SSL traffic within TCP streams whether standard or non-standard TCP ports are used. It is compatible with most protocols layered on SSL (for example, HTTP, SMTP, POP3, IMAP, and so on). The SSL appliance is also compatible with selected protocols which first send non-encrypted requests and responses, followed by the actual SSL protocol setup. The supported protocol variants that behave this way include the HTTP protocol's CONNECT method (used to traverse proxies) and the STARTTLS command (used by email protocols SMTP, POP3 and IMAP).

The SSL appliance can decrypt most SSL v3.0, TLS v1.0, TLS v1.1, and TLS v1.2 secured traffic (not just HTTPS traffic). The SSL appliance decrypts information received from the client, and re-encrypts it before sending it to the server, with the converse being performed for server to client traffic. Client and server software do not need to be modified, and security is maintained for the entire path between the client and the server.

SSL Decryption Methods

The SSL appliance supports two different methods that allow inspection of SSL. Each method requires that different information is available to the SSL appliance.

- *Known server key mechanism* relies on the SSL appliance having a copy of the SSL server's private key and certificate.
- *Certificate re-sign mechanism* relies on the SSL appliance having a trusted CA certificate that can be used to sign SSL server certificates that have been intercepted and modified.

Both methods can be used when the SSL appliance is operating in active-inline (see [Active-Inline Mode](#) on page 33) or passive-inline (see [Passive-Inline Mode](#) on page 32) mode but only the *known server key* method can be used if the SSL appliance is operating in passive-tap (see [Passive-Tap Mode](#) on page 29) mode.

IMPORTANT! It is possible for an SSL appliance to be configured to use both mechanisms at the same time, depending on details related to the SSL session.

There are different variations of these two basic mechanisms that are used depending on the type of SSL session being decrypted, the mode of operation of the SSL appliance, and the type of certificates/keys available to the system. The different variations are shown in detail in [Ruleset Policies](#) on page 37.

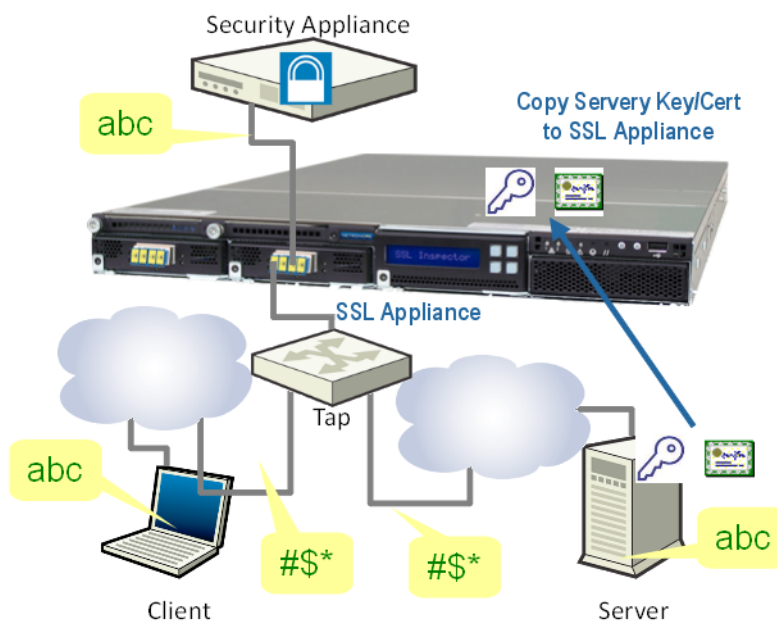
Known Server Key Method

This method (depicted below) illustrates the use of known server key decryption when the SSL appliance is connected in passive-tap mode. When the SSL appliance is deployed, the server certificate and key are installed on the SSL appliance for every server where you want to inspect traffic. The SSL appliance can use the key/certificate from a specific server to decrypt SSL sessions established with that server. A variant of this method requires that the server private key is installed only on the SSL appliance. If the private key only mode is being used, the references to key and certificate in the rest of this section should be taken to mean only the private key.

This method can be used only where the SSL appliance administrator has access to the server private key and certificate information. This is normally the case only if the SSL appliance and the server are managed and operated by the same organization or enterprise, that is, for *inbound* traffic to *your* servers.

The simplest example of known server key mode is shown in the following diagram:

Known Server Key Decryption Method -Passive-Tap Mode



This diagram shows that the client is sending *abc* to the server, which is encrypted to *\$\$** before being sent across the network. The server receives *\$\$** and decrypts it back to *abc* in order that the communication is successful. The SSL appliance receives a copy of the encrypted traffic *\$\$** from the tap device and, using the server key and certificate that have been loaded, decrypts this to get the plaintext *abc*.

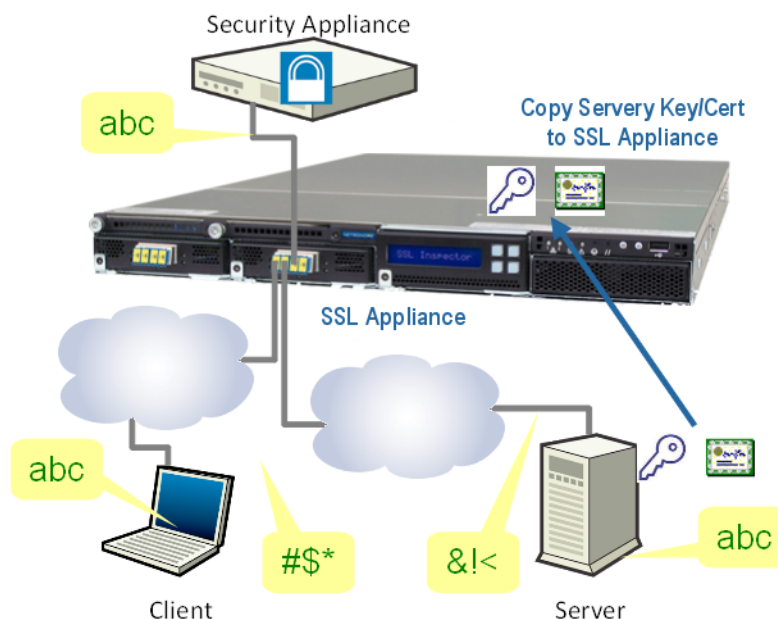
In this example, the SSL appliance is not a *man-in-the-middle* (MITM) of the SSL session it is simply receiving a copy of the encrypted data and decrypting it using the server private key and certificate that it has stored.

The fact that in passive-tap mode the SSL appliance is not a MITM for the SSL session is important because it means that not all SSL traffic can be decrypted even when the SSL appliance has the relevant servers private key and certificate. If the SSL session handshake makes use of Diffie-Hellman during the key exchange process then it is impossible for the SSL appliance to decrypt the traffic. To use known server key decryption to inspect a flow that uses Diffie-Hellman for key exchange the SSL appliance must be a MITM of the SSL session.

A maximum of 8192 known server key/certificate pairs can be loaded into the SSL appliance so traffic to many different SSL servers, with different SSL server certificates, can be inspected by a single SSL appliance.

The following diagram shows an example of known server key decryption when the SSL appliance is installed in passive-inline mode:

Known Server Key Decryption Method -Passive-Inline Mode



In this case, the SSL appliance is an MITM because the traffic between client and server passes through the SSL appliance. There are now two different encrypted SSL sessions. The client encrypts *abc* to *#\$** and sends this out over the network. Using its copy of the server private key and certificate, the SSL appliance decrypts this to access the plaintext *abc*. The SSL appliance re-encrypts the plaintext to produce *&!<* and sends this over the network to the server which can decrypt it to access the plaintext *abc*. The encrypted traffic

between the client and the SSL appliance, and between the SSL appliance and the server, is different. This is because there are two SSL sessions with different cryptographic session details in this example. If the session uses Diffie-Hellman for key exchange, the session details will be different for the two SSL sessions but if Diffie-Hellman is not used for key exchange, the session details can be the same and the SSL appliance can optimize performance by avoiding the need to re-encrypt the plaintext and simply forwarding the encrypted packet received from the client.

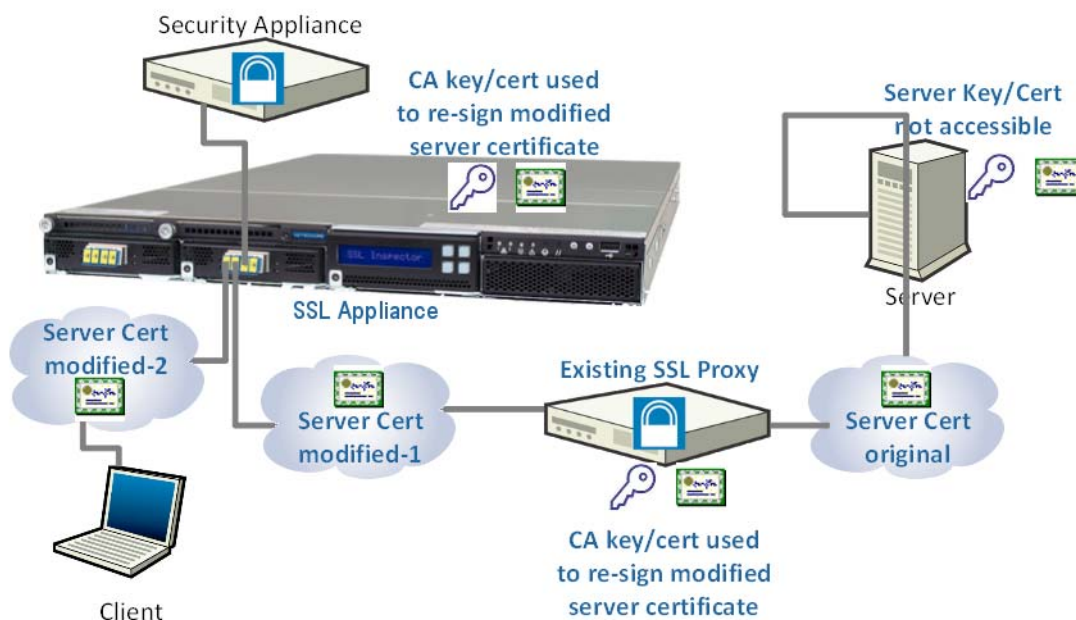
Certificate Re-Signing Method

Certificate re-sign is used when it is impossible to obtain a copy of the SSL server's private key and certificate. This is normally the case for any SSL servers not controlled by the organization deploying the SSL appliance. In general, any *outgoing* SSL traffic from an organization is inspected using certificate re-sign.

IMPORTANT! To use certificate re-sign, the SSL appliance must be an MITM, which means this mechanism cannot be used if the SSL appliance is connected in passive-tap mode.

The following diagram shows how certificate re-sign works:

Certificate Re-Sign Decryption Method -Passive-Inline Mode



The client initiates an SSL session to the server and the server responds by sending its SSL server certificate to the client. Because all traffic between client and server passes through the SSL appliance, it can detect and intercept the server certificate. After the SSL appliance intercepts the server certificate, it replaces the server's public keys with its own public keys and modifies the Certificate Revocation List (CRL) details in the server certificate. The SSL appliance then re-signs the server certificate using a Certificate Authority (CA) certificate and CA private key that is installed in the SSL appliance. The re-signed server certificate is then sent over the network to the client. As long as the client trusts the CA that was used to sign the server certificate it receives, it will not generate any warnings. Because the modified server certificate now contains public keys that are associated with private keys within the SSL appliance, it is possible for the SSL appliance to inspect the traffic.

When certificate re-sign is used, the two SSL sessions will always have different cryptographic session details and the SSL appliance must re-encrypt the plaintext before sending it back out to the network.

The client must trust the CA used to re-sign the server certificate. If the client does not trust the CA, it will generate warnings indicating that the SSL session should not be trusted. To ensure that the client does trust the CA used by the SSL appliance, there are two approaches which can be taken:

- The SSL appliance generates a CA certificate and keys internally and uses these to re-sign server certificates. The CA certificate which includes the CA public key is exported from the SSL appliance and then imported into the trusted CA store on the client; this is required only once.
- The SSL appliance is deployed in a network that has a private public key infrastructure (PKI) which is used to issue an intermediate CA certificate and keys that is loaded into the SSL appliance. Because the intermediate CA is issued by the enterprise root CA, it will automatically be trusted by all clients in the enterprise as will all server certificates that are signed by the intermediate CA.

Self-Signed Server Certificate Handling

Some SSL servers have server certificates that are self-signed, meaning the server generated the certificate and keys and then signed the certificate itself rather than having the certificate signed by a Certificate Authority (CA). Self-signed certificates are inherently less trustworthy than certificates signed by a trusted CA and some organizations may have a policy of not allowing SSL connections to servers that are using a self-signed certificate; the SSL appliance can be used to enforce such policies (see [Ruleset Policies](#) on page 37).

If SSL connections to servers using self-signed certificates are allowed, the SSL appliance can inspect the traffic using two different methods. The first method is to re-sign the certificate the same way a self-signed certificate is re-signed, see [Certificate Re-Signing Method](#) on page 23 for more details. This method is used if Decrypt (Resign) mode is chosen. The second method involves the self-signed

certificate information (for example, serial number, subject, and issuer) not being modified and only the public key and signature in the X.509 structure being replaced, effectively keeping the certificate self-signed. This method is used if **Replace Key Only** mode is chosen.

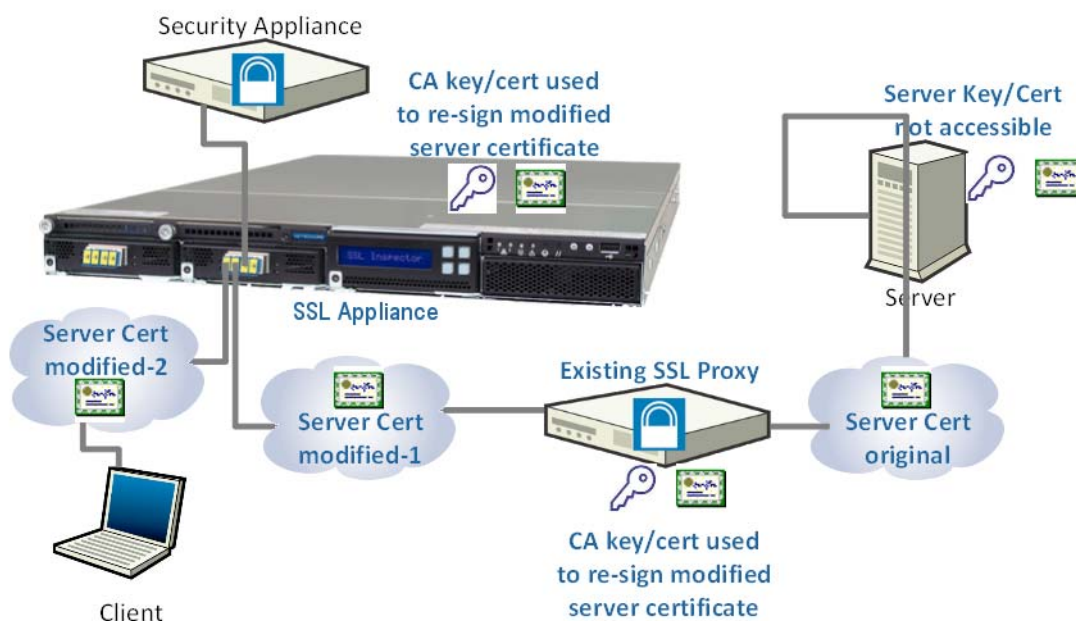
If the SSL appliance policy control has been used to block all traffic to servers using self-signed certificates, it is possible to explicitly allow traffic to a specific server using a self-signed certificate by loading a copy of the self-signed certificate into the Trusted Certificates store in the SSL appliance.

Decryption Methods in Cooperative Configurations

In some circumstances the SSL appliance may be deployed in networks that already have an SSL proxy device that is inspecting some of the outgoing SSL traffic using certificate re-sign. The SSL appliance would typically be deployed to allow other security appliances to view inspected traffic in addition to the existing proxy device that may not have an ability to pass inspected traffic to other devices. There are two possible ways to address this type of deployment and these are detailed below.

The following diagram shows a cooperative configuration with the SSL appliance deployed in passive-inline mode and using certificate re-sign:

Certificate Re-sign Decryption Method in a Cooperative Deployment

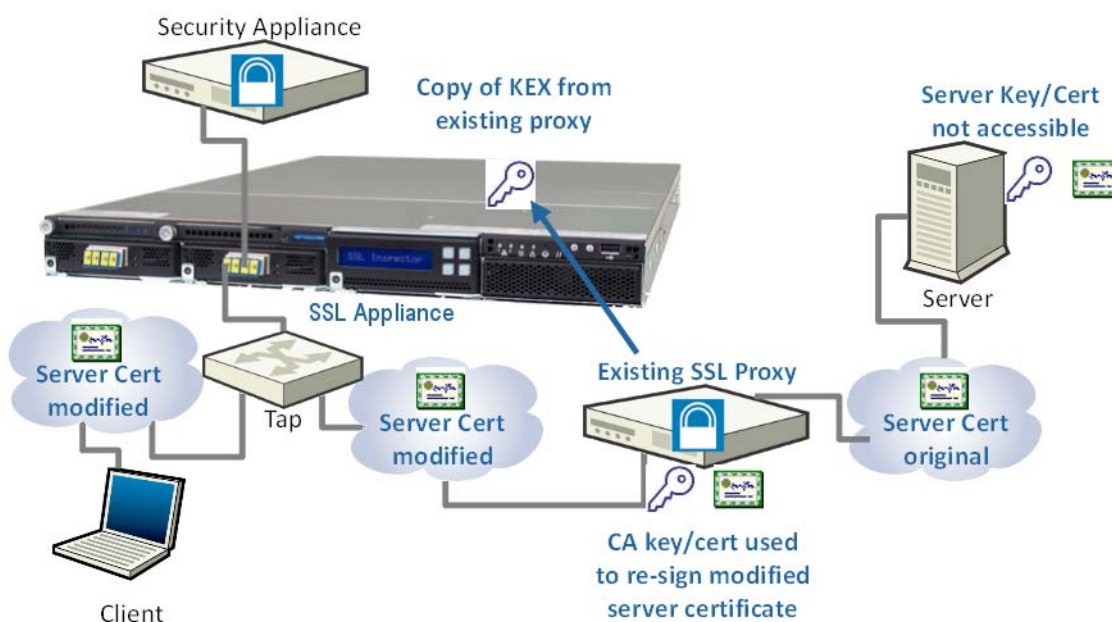


In this configuration both the existing SSL proxy and the SSL appliance are MITM devices. The existing proxy re-signs the original server certificate and then the

SSL appliance re-signs the modified server certificate it receives. In order for this configuration to work the SSL appliance must trust the CA that the existing proxy uses to re-sign server certificates and the client must trust the CA used by the SSL appliance. To simplify things, it is possible to add the CA used by the existing proxy to the trusted CA store in the SSL appliance and to use the same CA in the SSL appliance for certificate re-sign. This avoids the need for multiple CA certificates and removes the need to add an additional CA to the trust store on the client.

Another possible cooperative configuration is shown in the following diagram.

KEX Decryption Method in a Cooperative Deployment



This mechanism relies on being able to obtain the KEX key from the existing SSL proxy. The KEX is the key used by the existing proxy to encrypt/decrypt traffic. If the SSL appliance has a copy of the KEX, it is possible for it to inspect the session even if it is connected in passive-tap mode. Conceptually, this is similar to the known server key method of inspection though the details are not the same.

Marking SSL Plaintext

The generated flow containing plaintext obtained from inspected SSL traffic can be marked by the SSL appliance by modifying the source MAC address or by adding a VLAN tag to allow an attached device to distinguish this traffic from other traffic that was not inspected. In active-inline mode, a marking method must be selected because the SSL appliance must be able to distinguish returned plaintext traffic from other forwarded traffic. In passive-tap or passive-inline

mode, it is optional to have generated text marked. If modifying the source MAC address is enabled, the source MAC address is always set to 00: 15: 4D: 00: 00: D5. The VLAN tag value can be specified as part of the segment configuration.

Deployment Modes

This section explains how the SSL appliance is deployed in a network and how it operates in each of the deployment modes. Each segment is configured with a deployment mode that uses one or more network interfaces on the SSL appliance. There may be one or more segments configured on a single SSL appliance. Each segment is independent of the others segments. A network interface can only be associated with a single segment.

The following terminology is common to all deployment modes:

- Network port – a network interface that is either part of the bump-in-the-wire or is connected to a network tap device.
- Device port – a network interface that is connected to the primary attached appliance that is dealing with inspected traffic from the SSL appliance.
- Copy port – a network interface that is connected to a secondary passive appliance that is receiving a copy of the inspected traffic.
- Aggregation port – a network interface that is providing a connection to an additional network tap so that a segment can receive traffic from more than one network tap.
- Symmetric traffic – describes a situation where the packets for both directions of a network flow are seen on the same network interface on the SSL appliance.
- Asymmetric traffic – describes a situation where the packets for different directions of a network flow are seen on different network interfaces on the SSL appliance.
- Active-active – describes an HA deployment scenario where packets on a given flow may sent over either of the HA network links. From the SSL appliance's perspective this is equivalent to the Asymmetric traffic scenario in that packets belonging to a single flow may arrive on either one of two different network interfaces.

There are three main deployment modes for the SSL appliance with many variants within each mode. The following sections describe the way each of the


modes operates, for details on how to configure a segment and its mode of operation refer to the appropriate sections.

TIP! The actual physical interfaces on an SSL appliance that are used by a particular segment are allocated when the segment is activated, the webUI allows the user to choose the network interfaces to be used from the set of interfaces that are not currently in use by active segments.

When configuring an SSL appliance, consider the following questions:

- How many network interfaces connect traffic to the SSL appliance?
In a passive-tap mode there must be at least one network interface. In an inline mode the minimum number will be two as the SSL appliance is a bump-in-the-wire.
- Is the traffic being inspected symmetric or asymmetric?
If the traffic is asymmetric then more network interfaces will be required as the SSL appliance must see the packets for both directions of an SSL flow if it is going to be able to inspect the flow.
- Is there an active appliance connected to the SSL appliance?
An active appliance will require a minimum of two interfaces connecting it to the SSL appliance.
- Are there any passive appliances connected to the SSL appliance?
A passive appliance will require a minimum of one interface connecting it to the SSL appliance.
- Is there more than one passive appliance connected to the SSL appliance?
If more than one passive appliance is connected then should all traffic be copied to each passive appliance or should it be load balanced between the passive appliances.

The following example illustrates an active-inline deployment with one active appliance and load-balanced mirroring to two passive appliances.



- Arrows in and out on the left side are to and from the network being monitored (0 and 3).
- Arrows in and out on the right side are to and from the attached appliances (1,2,4 and 5).
- Ports shown in green are used by this segment.
- Ports shown in brown are not used by this segment.
- Ports below the dotted line (4 and 5) are used to mirror or load balance traffic from ports above the dotted line.
- Ports on the right with a dotted line between them (4 and 5) are used to load balance traffic across the two ports.
- Black arrows indicate the direction of packets in the flow on the indicated port.
- Fail-to-wire can connect pairs of ports together if activated (port pairs are 0&1, 2&3, 4&5, 6&7, and so on).

Deployment mode names have a standard structure illustrated by the following examples:

- `PT_outx1` indicates passive-tap with single output.
- `AI_FTA_1xcp2to1` indicates active-inline with fail-to-appliance and one mirrored output equivalent to a bi-directional tap.

Passive-Tap Mode

This section provides details on passive-tap modes of operation supported by the SSL appliance. Passive-tap mode connectivity options fall into three groups: symmetric, asymmetric, and aggregated.

- For symmetric traffic, the SSL appliance is connected to a single tap device that provides traffic for both directions of a flow over the single (bi-directional) tap port.
- For asymmetric traffic, the SSL appliance is connected to two tap devices with each tap device providing traffic for one direction of the flow.
- For aggregated traffic, the SSL appliance is connected to more than one bi-directional tap port and aggregating traffic from all the tap ports into a single segment.

IMPORTANT! Only known server key decryption and KEX decryption methods can be used when the SSL appliance is deployed in passive-tap mode.


IMPORTANT! When Diffie-Hellman is used for key exchange, the SSL appliance is unable to decrypt the flow using known server key or KEX methods when connected in passive-tap mode.

An SSL appliance in passive-tap mode may be configured to log sessions but not inspect SSL traffic. This configuration collects session log data on all of the SSL traffic in the network and does not require access to any certificates or keys. Analysis of the session log provides a detailed picture of the SSL traffic in the network and can be used to plan what traffic must be inspected and how the SSL appliance will connect to the network to achieve this.


The simplest passive-tap modes deal with symmetric traffic being inspected.

The simplest passive-tap deployment has the SSL appliance connected to a tap that delivers symmetric traffic to the SSL appliance over a single network interface. The inspected traffic is then sent to a single passive appliance as symmetric traffic over a single network interface.

TIP! If two tap ports are being used in aggregation mode and are connected to interfaces that share fail-to-wire hardware then whenever the fail-to-wire is active the two taps will be connected to each other. You must ensure that this will not cause problems for the tap ports or the network.




The following deployments use the aggregation capabilities of the SSL appliance to combine traffic from two or three network taps onto a single SSL appliance segment. In both these examples the inspected traffic is sent to a single attached appliance as symmetric traffic over a single interface (Device port).




Any of the above modes can use an additional two interfaces (copy ports) to connect to additional attached passive appliances. If only one copy port is configured then it will feed a copy of the symmetric traffic from the SSL appliance to a second passive appliance. If two copy ports are used then these can be used to:

- Feed a copy of the symmetric traffic to a second and third passive appliance
- Feed an asymmetric copy of the traffic to a second passive appliance
- Load balance the symmetric traffic to a second and third passive appliance

The copy options for all three of the above operating modes are shown below.



Passive-tap mode that supports inspection of asymmetric traffic is shown below. The figure on the right shows the copy options available for this mode of operation.



The first four modes are for an SSL appliance connected to a single bi-directional tap port that provides traffic for both directions of a flow. They differ in terms of how non-SSL and inspected SSL traffic is sent to the attached passive security appliances.

If no copy ports are used then a single passive appliance will receive the asymmetric traffic from the SSL appliance over the two device ports.

If a single copy port is used then it will feed a symmetric copy of the asymmetric traffic from the SSL appliance to a second passive appliance. If two interfaces are used then you can:

- Feed a copy of the asymmetric traffic to a second passive appliance
- Feed a symmetric copy of the traffic to a second and third passive appliance
- Load balance the symmetric traffic to a second and third passive appliance

If four interfaces are used then you can:

- Feed a copy of the asymmetric traffic to a second and third passive appliance
- Load balance the asymmetric traffic to a second and third passive appliance

These four diagrams show configurations with traffic into the SSL appliance from two span ports, one for each direction of traffic on a flow.


Passive-Inline Mode

This section provides details on all the different Passive-Inline modes of operation supported by the SSL appliance. Passive-Inline mode connectivity can be either symmetric or asymmetric:

- For symmetric traffic, the SSL appliance is connected inline on a network segment that carries traffic for both directions of a flow.
- For asymmetric traffic, the SSL appliance is connected inline on two network segments with packets for a given flow potentially being present on one or other segment.

TIP! If the SSL appliance is being deployed in a network using an active-active HA architecture then this can be treated as an asymmetric traffic case. The SSL appliance can be configured as an inline device in both active links in the HA network and will treat these as a single Segment internally. It does not matter which packets on a given flow occur on which of the active-active links.

The figures below show the simple passive-Inline configuration on the left and the copy port options that are available on the right. In passive-Inline mode there are no device ports configured as part of the initial segment configuration so all attached appliances are connected to copy ports.



If a single copy port is used then it will feed a symmetric copy of the symmetric traffic from the SSL appliance to the first passive appliance. If two copy ports are used then these can be used to:


- Feed a copy of the symmetric traffic to the first and second passive appliances
- Feed an asymmetric copy of the traffic to the first passive appliance
- Load balance the symmetric traffic to the first and second passive appliances

If four copy ports are used then these can be used to:

- Feed an asymmetric copy of the traffic to the first and second passive appliances
- Load balance an asymmetric copy of the traffic to the first and second passive appliances

Passive-Inline mode that allows inspection of asymmetric traffic is shown below on the left and the copy port options available are shown in below on the right. In

passive-Inline mode there are no device ports configured as part of the initial segment configuration so all attached appliances are connected to copy ports.



If a single copy port is used then it will feed a symmetric copy of the symmetric traffic from the SSL appliance to the first passive appliance. If two copy ports are used then these can be used to:

- Feed a copy of the symmetric traffic to the first and second passive appliances
- Feed an asymmetric copy of the traffic to the first passive appliance
- Load balance the symmetric traffic to the first and second passive appliances

If four copy ports are used then these can be used to:

- Feed an asymmetric copy of the traffic to the first and second passive appliances
- Load balance an asymmetric copy of the traffic to the first and second passive appliances

Active-Inline Mode

This section provides details on all the different Active-Inline modes of operation supported by the SSL appliance. Active-Inline mode connectivity can be symmetric or asymmetric:


- For symmetric traffic, the SSL appliance is connected inline on a network segment that carries traffic for both directions of a flow.
- For asymmetric traffic, the SSL appliance is connected inline on two network segments with packets for a given flow potentially present on one or the other segment.

TIP! If the SSL appliance is deployed in a network using an active-active HA architecture then this can be treated as an asymmetric traffic case. The SSL appliance can be configured as an inline device in both active links in the HA network and will treat these as a single segment internally. It does not matter which packets on a given flow occur on which of the active-active links.

All Active-inline modes of operation have an active appliance attached to the SSL appliance via the device ports, the way in which the active appliance is connected determines how traffic flows in the event of a failure of the SSL

appliance. Fail To Appliance (FTA) mode results in traffic flowing through the attached active appliance in the event of failure while Fail To Network (FTN) mode results in traffic bypassing the active appliance in the event of failure.

The diagrams below show Active-inline modes for situations where symmetric traffic is passing through the SSL appliance. The diagram on the right shows the copy port options available in Active-inline mode.




If a single copy port is used then it will feed a symmetric copy of the symmetric traffic from the SSL appliance to the first passive appliance. If two copy ports are used then these can be used to:

- Feed an a symmetric copy of the traffic to the first passive appliance
- Load balance the symmetric traffic to the first and second passive appliances

If four copy ports are used then these can be used to:

- Feed an asymmetric copy of the traffic to the first and second passive appliances
- Load balance an asymmetric copy of the traffic to the first and second passive appliances

Active-inline mode for dealing with asymmetric traffic is shown in the diagrams below.



If a single copy port is used then it will feed a symmetric copy of the asymmetric traffic from the SSL appliance to the first passive appliance. If two copy ports are used then these can be used to:

- Feed an a symmetric copy of the traffic to the first passive appliance
- Load balance the symmetric traffic to the first and second passive appliances

If four copy ports are used then these can be used to:

- Feed an asymmetric copy of the traffic to the first and second passive appliances
- Load balance an asymmetric copy of the traffic to the first and second passive appliances

Policies

Policies in the SSL appliance are composed of three elements:

- List
- Segment
- Ruleset

A list is used to collect multiple items of the same type of information so that a single policy can point to the list and the policy is applied whenever any of the items in the list are true. For example, a list may contain twenty different Common Names (CN) that occur in the server certificates from twenty different sites. A policy that is configured to inspect traffic when it detects a particular CN can point to the list instead of just indicating a single CN in the policy. This allows a single policy entry to apply to all twenty different sites and means that additional sites can be added (by editing the list) without needing to edit the policy.

A segment contains some policy information and is linked to a *ruleset* that contains the majority of the policy information. Lists are used within rulesets to make it easier to have policies that apply to many different SSL sessions. The system can have multiple segments defined and can have more than one segment active at any point in time. For example, a system could have six rulesets defined (*ruleset1* to *ruleset6*) and might have two active segments each using different ports on the SSL appliance. Segment 1 can use *ruleset1* and segment 2 can use *ruleset4* or both segments can use *ruleset3*. Inactive segments are not associated with physical ports on the SSL appliance until the point at which they are activated.

A segment is created by selecting one of the deployment modes, described in [Deployment Modes](#) on page 27. The system will then allocate external ports on the SSL appliance that are used by this segment when it is activated. As part of creating the segment, a number of default policy actions are defined which apply specifically to the segment. Some of these can be overridden by more explicit policies that are defined in the ruleset associated with this segment.

Policies can be used in the SSL appliance to control the following:

- Which SSL sessions are inspected
- What decryption method is used to inspect a specific session
- Whether an SSL session that is not being inspected is cut through or dropped

- Whether SSL sessions using specific cipher suites are allowed across the network
- How SSL sessions that cannot be decrypted are handled
- How SSL sessions with specific certificate status are handled
- How SSL sessions to servers using self-signed certificates are handled

Segment Policies

The policies that form part of the segment definition are created with default values which can then be modified. A segment contains policy settings as shown in the [Segment Policy Options](#) table.

Segment Policy Options

Item	Default Setting	Notes
Name		Identifies this segment configuration
Comment		Optional descriptive text
Mode		Operating mode for segment chosen from list
Ruleset		Name of ruleset used by segment
Session Log	Disabled	Enable or disable SSL session log for this segment
Compression	Cut-through	This block has policy definitions for how SSL flows that cannot be decrypted are handled on this segment. The cipher suite setting consults a list of cipher suites that cannot be decrypted by the SSL appliance.
SSL v2	Cut-through	
TLS v1.2	Cut-through	
Diffie-Hellman passive-tap mode	Cut-through	
Client Certificate	Reject	
Cipher suite	Cut-through	
Uncached session	Cut-through	

Segment Policy Options (Continued)

Item	Default Setting	Notes
Invalid Issuer		This block has policy definitions that define how to handle specific conditions that occur in the SSL server certificate for a session. The Segment/Rule priority setting determines whether a rule in the ruleset takes priority or is overridden by the segment rule.
Invalid Signature		
Expired		
Not yet valid		
Self-signed		
Segment/rule priority	Rule over Segment	

Ruleset Policies

A ruleset contains a basic set of defined items and a variable number of rules. A rule can be specific to a given type of SSL flow or can apply to multiple SSL flows that are defined in a list. For example, a list may contain the Common Names (CN) of ten SSL servers and a rule that references this list will be applied to any SSL sessions where the CN in the server certificate matches anyone of the ten names in the list. In the following tables, any entry where the Default Setting field is empty means that the default setting is the **nothing is set** option.

CN and DN entries, whether in lists or as single entries, can include a * wildcard at the beginning of the entry. This allows for one entry to match on multiple CNs or DNs. For example a CN value of *.googl e. com will match any SSL session with a CN that ends in .googl e. com.

The [Ruleset Policy Options](#) table shows the basic set of policy options contained in a ruleset. Note that there will typically be many rule items in a single ruleset. The details relating to rules themselves are shown in more detail later in this section.

Ruleset Policy Options

Item	Default Setting	Notes
Name		Identifies this ruleset
Default Internal Certificate Authority		Default CA used for certificate re-sign
External Certificate Authorities	All external CAs	Can point to a custom list instead

Ruleset Policy Options (Continued)

Item	Default Setting	Notes
Certificate Revocation Lists	All CRL lists	Can point to a custom list instead
Trusted Certificates		Optional list
Catch All Action	Cut through	Catch all action – cut, reject, or drop
Rules		Rules are of different types (see below) depending on what action they specify

There are seven types of rules that can occur within a ruleset, and any type can occur multiple times or not at all in a given ruleset. Each rule contains multiple match fields that can be configured. These fields are compared with the corresponding values in an SSL session to determine whether the rule should be applied to the session. Any match fields that are left empty are treated as matching any value for that field. The seven rule types allow for a total of nine possible actions that can be taken if a rule is matched. These actions are listed in the [Rule Actions](#) table.

Rule Actions

Action	Type ID
Decrypt (Certificate and Key known)	1
Decrypt (Key known)	2
Replace Key Only	3
Replace Certificate and Key	4
Decrypt (Resign Certificate)	5
Decrypt (Anonymous Diffie-Hellman)	6
Cut-Through	7
Drop	7
Reject	7

Note that some of the match fields can point to lists which allows a single rule entry to be triggered by more than one set of matching criteria. If there is a field

which points to a specific item and another field which points to a list of these items, the fields are mutually exclusive; only one of the fields can be used. In the following tables, mutually exclusive fields are shown by arrows (↑↓) in the default setting column.

If a rule in a ruleset cannot be applied due to the mode of operation of the segment, it will be ignored and a warning will be logged. For example, a rule that specifies decryption using certificate re-sign cannot be applied if the segment is operating in passive-tap mode.

The [Decrypt with Known Certificate and Key Rule Format](#) table shows details for a Decrypt (Certificate and Key known) rule that will trigger decryption using a known server key and certificate if the details in the server certificate for a session match the rule.

Decrypt with Known Certificate and Key Rule Format

Item	Default Setting	Notes
Decrypt (Certificate and Key known)		Decrypt using known key and certificate
Comment		Optional descriptive text
Known Certificate with Key	↓	Pointer to a single certificate/key value
Known Certificates with Keys	↑ All Known	Name of a list of certificate/key pairs that is checked for a match
Source IP	↓	IP address and subnet mask
Source IP List	↑	Name of list of source addresses/masks that is checked for a match
Destination IP	↓	IP address and subnet mask
Destination IP List	↑	Name of list of destination addresses/masks that is checked for a match
Destination Port		Destination IP port number

The [Decrypt with Known Key Rule Format](#) table shows details for a Decrypt (Key known) rule that will trigger decryption using a known key decryption method if the details in the server certificate for a session match the rule. Note that some of

the match fields can point to lists which allows a single rule entry to be triggered by more than one set of matching criteria.

Decrypt with Known Key Rule Format

Item	Default Setting	Notes
Decrypt (Key known)		Decrypt using known key
Comment		Optional descriptive text
Known Key	↓	Pointer to a single known key value
Known Keys	↑	Name of a list of known keys that is checked for a match
Source IP	↓	IP address and subnet mask
Source IP List	↑	Name of list of source addresses/masks that is checked for a match
Destination IP	↓	IP address and subnet mask
Destination IP List	↑	Name of list of destination addresses/masks that is checked for a match
Destination Port		Destination IP port number
Certificate Status		Status of X.509 server certificate

The [Decrypt Using Key Replacement Format](#) table shows details for a Replace Key Only rule that will trigger decryption using a certificate modification method for a self-signed certificate if the details in the self-signed server certificate for a session match the rule. Note that some of the match fields can point to lists which allows a single rule entry to be triggered by more than one set of matching criteria.

Decrypt Using Key Replacement Format

Item	Default Setting	Notes
Replace Key Only		Decrypt using key replacement and randomize S/N
Comment		Optional descriptive text
Cipher Suite List		List of cipher suites – cannot include Anonymous Diffie-Hellman cipher suites

Decrypt Using Key Replacement Format (Continued)

Item	Default Setting	Notes
Trusted Certificate	↓	Trusted Certificate that is checked for a match
Trusted Certificates	↑	List of Trusted Certificates that are checked for a match
Subject DN	↓	X.509 Subject Distinguished Name that is checked for a match
Subject DN List	↑	List of X.509 Subject Distinguished Names that are checked for a match
Issuer DN	↓	X.509 Issuer Distinguished Name that is checked for a match
Issuer DN List	↑	List of X.509 Issuer Distinguished Names that are checked for a match
Source IP	↓	IP address and subnet mask
Source IP List	↑	Name of list of source addresses/masks that is checked for a match
Destination IP	↓	IP address and subnet mask
Destination IP List	↑	Name of list of destination addresses/masks that is checked for a match
Destination Port		Destination IP port number

The [Decrypt Using Replacement of Key and Certificate Format](#) table shows details for a Replace Certificate and Key rule that will trigger decryption using a certificate and key replacement method if the details in the server certificate for a session match the rule. Note that some of the match fields can point to lists

which allows a single rule entry to be triggered by more than one set of matching criteria.

Decrypt Using Replacement of Key and Certificate Format

Item	Default Setting	Notes
Replace Certificate and Key		Decrypt using key and certificate replacement
Comment		Optional descriptive text
Known Certificate with Key (to replace with)		Pointer to a certificate and key that will be used to replace the certificate and key in the server certificate
Cipher Suite List		List of cipher suites – cannot include Anonymous Diffie-Hellman cipher suites
Trusted Certificate	↓	Trusted Certificate that is checked for a match
Trusted Certificates	↑	List of Trusted Certificates that are checked for a match
Subject DN	↓	X.509 Subject Distinguished Name that is checked for a match
Subject DN List	↑	List of X.509 Subject Distinguished Names that are checked for a match
Issuer DN	↓	X.509 Issuer Distinguished Name that is checked for a match
Issuer DN List	↑	List of X.509 Issuer Distinguished Names that are checked for a match
Source IP	↓	IP addresses and subnet masks
Source IP List	↑	List of source addresses/masks that is checked for a match
Destination IP	↓	IP addresses and subnet masks
Destination IP List	↑	List of destination addresses/masks that is checked for a match
Destination Port		Destination IP port number
Certificate Status		Status of X.509 server certificate

The [Decrypt Using Certificate Re-sign Format](#) table shows details for a Decrypt (Re-sign Certificate) rule that will trigger decryption using certificate re-sign if the details in the server certificate for a session match the rule. Note that some of the match fields can point to lists which allows a single rule entry to be triggered by more than one set of matching criteria.

Decrypt Using Certificate Re-sign Format

Item	Default Setting	Notes
Decrypt (Re-sign Certificate)		Decrypt using certificate re-sign
Comment		Optional descriptive text
Internal CA		Pointer to the internal CA that is used to re-sign the server certificate
Cipher Suite List		List of cipher suites – cannot include Anonymous Diffie-Hellman cipher suites
Trusted Certificate	↓	Trusted Certificate that is checked for a match
Trusted Certificates	↑	List of Trusted Certificates that are checked for a match
Subject DN	↓	X.509 Subject Distinguished Name that is checked for a match
Subject DN List	↑	List of X.509 Subject Distinguished Names that are checked for a match
Issuer DN	↓	X.509 Issuer Distinguished Name that is checked for a match
Issuer DN List	↑	List of X.509 Issuer Distinguished Names that are checked for a match
Source IP	↓	IP address and subnet mask
Source IP List	↑	Name of list of source addresses/masks that is checked for a match
Destination IP	↓	IP address and subnet mask

Decrypt Using Certificate Re-sign Format (Continued)

Item	Default Setting	Notes
Destination IP List	↑	Name of list of destination addresses/masks that is checked for a match
Destination Port		Destination IP port number
Certificate Status		Status of X.509 server certificate

The [Decrypt Anonymous Diffie-Hellman Format](#) table shows details for a Decrypt (Anonymous Diffie-Hellman) rule that will trigger decryption using certificate re-sign if the details in the server certificate for a session match the rule. Note that some of the match fields can point to lists which allows a single rule entry to be triggered by more than one set of matching criteria.

Decrypt Anonymous Diffie-Hellman Format

Item	Default Setting	Notes
Decrypt (Anonymous Diffie-Hellman)		Decrypt Anonymous Diffie-Hellman session
Comment		Optional descriptive text
Source IP	↓	IP address and subnet mask
Source IP List	↑	Name of list of source addresses/masks that is checked for a match
Destination IP	↓	IP address and subnet mask
Destination IP List	↑	Name of list of destination addresses/masks that is checked for a match
Destination Port		Destination IP port number
Certificate Status		Status of X.509 server certificate

The [Rules That Do Not Involve Decryption Format](#) table shows details for Cut Through/Drop/Reject rules that will trigger actions other than decryption (for example, rules that cut sessions through, reject sessions or drop them if the details in the server certificate for a session match the rule). Note that some of

the match fields can point to lists which allows a single rule entry to be triggered by more than one set of matching criteria.

Rules That Do Not Involve Decryption Format

Item	Default Setting	Notes
Cut-Through/Drop/Reject		Actions are cut, reject or drop
Comment		Optional descriptive text
Cipher Suite List		List of cipher suites – cannot include Anonymous Diffie-Hellman cipher suites
Trusted Certificate	↓	Trusted Certificate that is checked for a match
Trusted Certificates	↑	List of Trusted Certificates that are checked for a match
Subject DN	↓	X.509 Subject Distinguished Name that is checked for a match
Subject DN List	↑	List of X.509 Subject Distinguished Names that are checked for a match
Issuer DN	↓	X.509 Issuer Distinguished Name that is checked for a match
Issuer DN List	↑	List of X.509 Issuer Distinguished Names that are checked for a match
Source IP	↓	IP address and subnet mask
Source IP List	↑	Name of list of source addresses/masks that is checked for a match
Destination IP	↓	IP address and subnet mask
Destination IP List	↑	Name of list of destination addresses/masks that is checked for a match
Destination Port		Destination IP port number
Certificate Status		Status of X.509 server certificate

Lists

Lists can be referenced by rules in rulesets and allow a single rule to be applied to more than one flow because any flow that matches an entry in the list will trigger the rule action. For each type of PKI list, the system will create a default list that is read only and includes all items of that type present in the system. The default list has names that begin with `all-`. User created custom lists are subsets of the default lists.

The [List Types and Contents](#) table shows the default set of lists that exist within the SSL appliance.

List Types and Contents

Name	Contains
<code>all-external-certificate-authorities</code>	All trusted external CAs
<code>all-certificate-revocation-lists</code>	All pointers to CRL locations
<code>all-known-certificates</code>	All known server certificates
<code>all-known-keys</code>	All known server private keys
<code>all-known-certificates-with-keys</code>	All known server private key/certificates
<code>ssl-ng-unsupported-sites</code>	Sites it is not possible to inspect SSL sessions to

Importing of new keys or certificates is always done to the relevant `all` list. Adding entries to a custom list is done by selecting entries from the relevant `all` list.

In addition to the above lists of PKI items the system can contain lists of:

- Distinguished Names
- Common Names
- Cipher Suites
- IP addresses

Reset Generation

There are several conditions under which the SSL appliance prematurely terminates TCP connections that pass through it using TCP RST packets. Presently, all of these conditions only apply when the SSL appliance is deployed in active-inline or passive-inline mode. The device does not terminate connections

prematurely in passive-tap mode. The appliance generates TCP RST packets when it receives a packet for a flow that triggers a Reset rule, when an undecryptable policy is triggered, or when there is an error in a flow that has been modified so that the remainder of the flow cannot be cut through.

When the SSL appliance determines that it must reject a TCP flow, it releases most of the state associated with that flow and considers the flow terminated. From that point on, the appliance will turn around any packets that it receives and determines to be a part of the original flow into RST packets and transmit them back to the sender. Thus, if any of the RST packets are lost, packets from the original client or server will trigger RSTs to hang up the connection. An administrator may configure the policy of the appliance to always reject certain flows whenever they arrive. In such a case, the SSL appliance will generate RSTs by turning around packets in flows matching the policy's pattern, but will not spontaneously generate RSTs to send to connection endpoints.

If the SSL appliance rejects a flow, the appliance also tries to signal both endpoints of the connection about the termination by generating a *spontaneous* TCP RST for each endpoint of the connection. After the initial rejection, any subsequently received packets for the same flow will continue to trigger RSTs back to the sender as described above. There is one special case for a flow rejection triggered by a TCP SYN. In such a case, there is no server endpoint or state; the SSL appliance generates only one spontaneous RST to send back to the SYN packet's source. Events that will cause the SSL appliance to generate RST packets are:

- Flows rejected because of an action configured for dealing with undecryptable flows (for example the presence of a client certificate in a flow that prevents it being inspected).
- Modified flows with decryption errors (where decrypt and re-encrypt are being done). As the flow is modified it cannot simply be cut through after the error.

If the SSL appliance is operating in active-inline mode then the attached inline appliance can also cause the SSL appliance to generate a reset in both directions on an SSL flow that is being inspected. If the inline appliance drops a packet from the generated TCP flow that is carrying the decrypted payload data then the SSL appliance will detect this and generate an RST in both directions on the original SSL flow in order to kill the flow. If the active appliance generates an RST itself on the generated TCP flow then this will be detected by the SSL appliance and will trigger an RST in each direction on the original SSL flow.

Failure Modes and High Availability

The SSL appliance can automatically respond to certain types of failures that it detects. The term *failure option* refers to a set of responses that the SSL appliance performs when it detects a particular type of failure.

There are two types of failures that the SSL appliance can detect and respond to:

- Link failure (interface going down); associated with a segment
- Software failure (data-plane); associated with the device

A segment can be configured to operate in normal mode or High Availability (HA) mode. The failure actions taken by the device differ depending on whether the segment is configured for HA mode or not. Because HA mode is not relevant if a segment is operating in passive-tap mode, HA mode can be configured only for segments operating in active-inline or passive-inline mode. The behavior in response to a link failure differs if a segment is operating in HA mode.

In High Availability (HA) mode the failure options are set up to enable the SSL appliance to propagate failure state to the ethernet switches that it is connected to so that the switches can direct traffic to an alternate SSL appliance system to maintain availability. When not in HA mode link state is not propagated between links on a segment.



Within the system, software failures are handled by a failure mode state machine while link failures are handled by a failure mode filter which is located before the failure mode state machine. If a segment is operating in HA mode, the failure mode filter is active; otherwise, it is disabled.

The following sections detail how link failures and software failures are dealt with and how segments can be configured to respond to the impact of such failures.

Link Failures



The effect of a link failure on a segment is not configurable, however the segment behavior is different depending on whether the segment is operating in HA mode. Configuring HA mode enables the failure mode filter which is otherwise inactive.

When not operating in HA mode, the failure of a link used only by the segment has the following impact:

- The link state for the affected link will go to down.
- The link status LEDs for the affected link will show that the link is down.
- The dashboard Network Interfaces status display will show the affected link as down.
- The dashboard Segments Status display will show the segment with a yellow background.
- The System status indicator will change to  in the status bar at the bottom of the screen.
- The Network status indicator will change to  in the status bar at the bottom of the screen.
- The event will be logged in the system log.

- Detection and inspection of SSL traffic will cease when the link is part of the bump-in-the-wire for an inline segment or is the link to the network tap in PT mode.
- Even though at least one of the attached passive appliances is no longer receiving the inspected traffic when the link is a link to an attached passive appliance, SSL detection and inspection will continue.

If the segment is operating in HA mode, the following actions will take place if a link being used by the segment goes down:

- Failure of any segment interface will force all the network facing interfaces in the segment down when the segment is passive-inline.
- When the segment is active-inline, failure of any segment interface, other than those used for mirroring, will force all non-mirrored interfaces in the segment down.
- The link state for the affected links will go to down.
- The link status LEDs for the affected links will show that the link is down.
- The dashboard Network Interfaces status display will show the affected links as down.
- The dashboard Segments Status display will show the segment with a red background.
- The System status indicator will change to  in the status bar at the bottom of the screen.
- The Network status indicator will change to  in the status bar at the bottom of the screen.
- The event will be logged in the system log.
- Detection and inspection of SSL traffic will cease.
- All data-plane failures will be ignored while a segment is in link failure mode.
- Recovery from link failure mode requires manual intervention (forced recovery from the WebUI).

Software (Data-Plane) Failures

Software failures are triggered by one or more checks that are run in the background while the device is operating. These background checks are for the system, not for a specific segment. The subsystem running the checks provides a keep alive watchdog signal to the failure engine. If the failure engine does not receive the keep alive indication, it triggers the failure mechanism.

The failure mode that becomes active when a failure occurs is configured per segment. A failure may trigger different failure modes for different segments if they are configured differently. Some of the failure modes require manual intervention to exit the mode while others will automatically exit as soon as the condition that caused the failure and any other failure conditions are removed. See [Segments](#) on page 119 for more details.

The various failure modes that can be configured for a segment are:

- Disable interfaces
- Drop packets (Auto Recovery)
- Fail-to-wire (Auto Recovery)
- Fail-to-wire (Manual Reset)
- Ignore failure

The options for High Availability mode are:

- Disabled – HA mode is not active
- Auto Recovery – automatic recovery from failure mode when the cause of the failure is removed
- Manual Reset – manual action via the webUI is needed to exit failure mode.

Modes that invoke fail-to-wire cause the hardware mechanisms in the NetMod to activate and connect together pairs of external ports to ensure that traffic continues to flow through the network while the SSL appliance is failed.

During a software failure state, any link state changes are processed because link failures have priority over software failures.

Internally the system generates a recovery event after the issues that caused the software failure have been removed and all run-time tests have succeeded. Automatic recovery will occur after the recovery event occurs as long as the segment is configured to use one of the automatic recovery modes. If a manual recovery mode is in operation, the manual reset will only be accepted after the system has generated a recovery event. Manual recovery is achieved by clicking on the **Manually Unfail** button on the dashboard. This button will be enabled only if Manual Unfail is allowed and if the condition which triggered the failure is not resolved.

Example Deployment Configurations

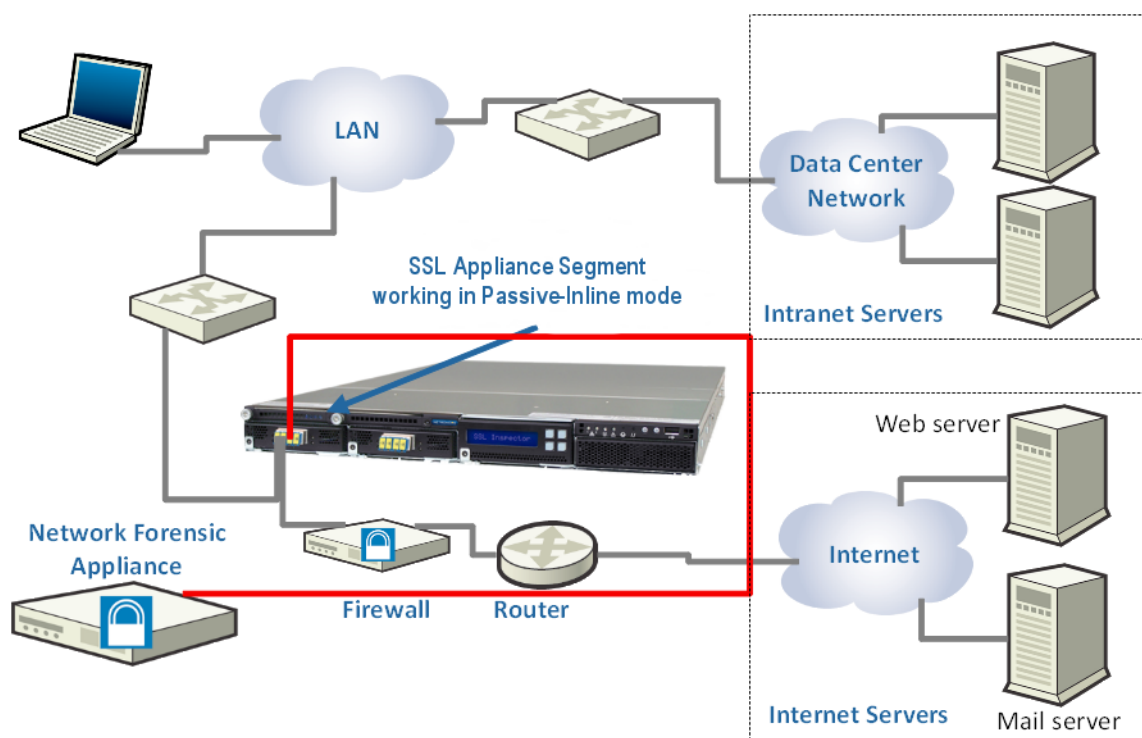
This section provides some examples of how the SSL appliance can be deployed alongside other security appliances to protect the network against threats carried by SSL traffic. Network links shown in red indicate links that are carrying decrypted SSL traffic.

Outbound Inspection

The following diagram shows an outbound monitoring scenario. The monitored web browsers or other SSL clients are located in the private network (intranet),

with the monitored servers typically being located in the Internet or in partner's extranets.

Outbound Monitoring with Network Forensic Appliance

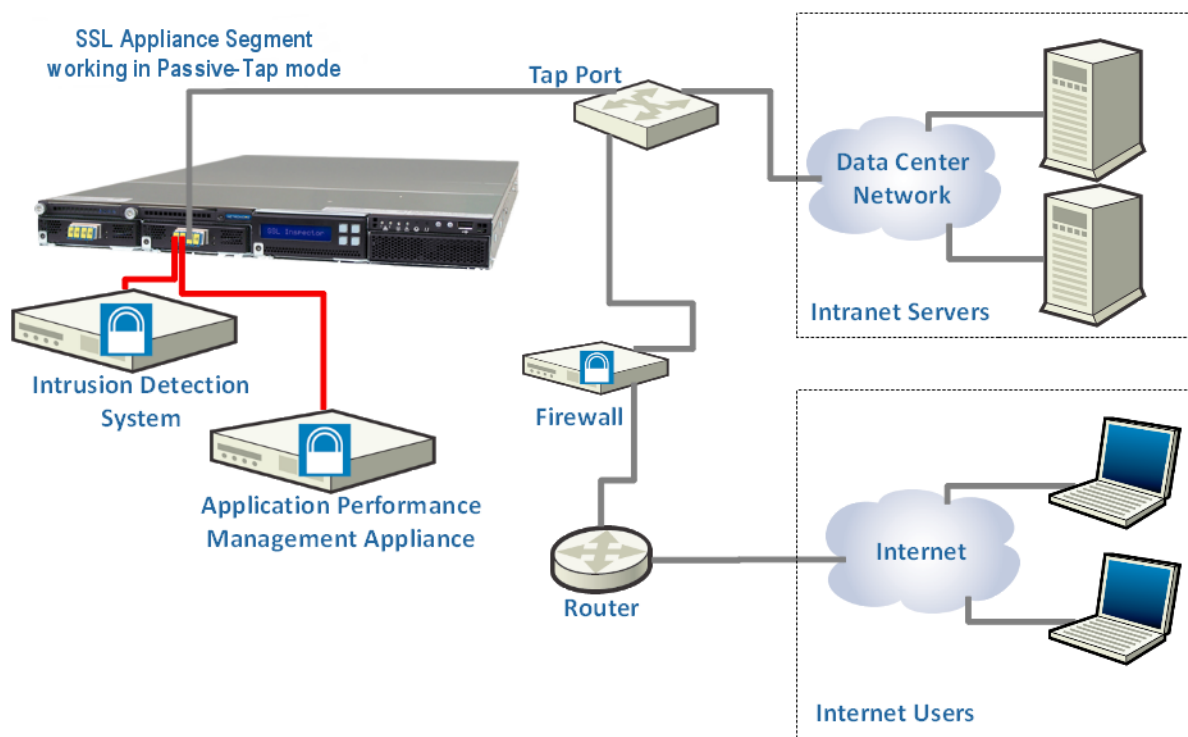


For this scenario the SSL appliance is typically deployed adjacent to the firewall or router which leads to the Internet. The SSL appliance must be deployed on the public side of the firewall if the firewall itself generates SSL-encrypted traffic that must be inspected (for example, if the firewall also includes SSL VPN capabilities) or if the network topology requires deploying the SSL appliance at that location (for example, because the firewall also aggregates multiple network segments). For all other cases, deploying the SSL appliance on the private side of the firewall is advisable. In this deployment, traffic is inspected using certificate re-sign (see [Certificate Re-Signing Method](#) on page 23) because the SSL servers are not under the control of the enterprise deploying the SSL appliance and it is not possible to obtain copies of the server private key or certificate for these servers. The client systems in this deployment must trust the Certificate Authority used by the SSL appliance to re-sign server certificates. The [PI_1xcp2t01](#) diagram in [Passive-Inline Mode](#) on page 32 shows the connection mode being used in this example.

Inbound Inspection

The following diagram shows a deployment where the SSL appliance is connected to a network tap or span port and is delivering decrypted traffic to an Intrusion Detection System and to an Application Performance Monitoring system.

Inbound Monitoring with IDS and Application Performance Monitor



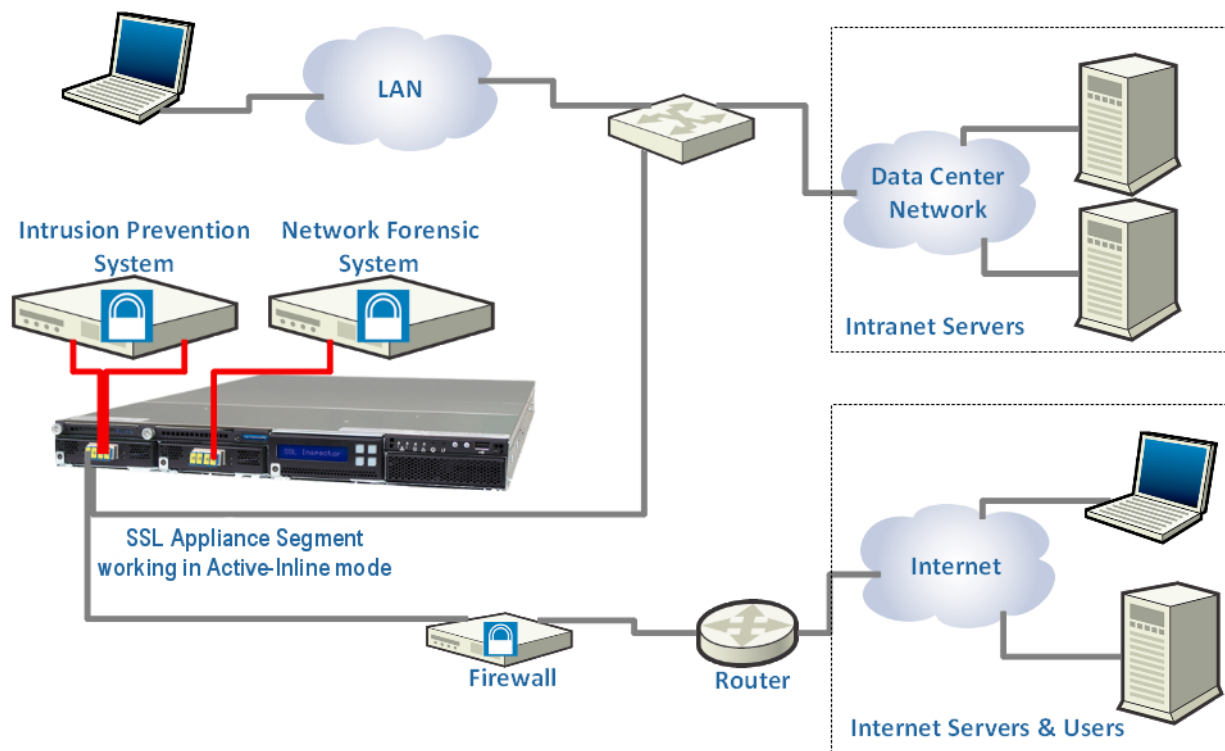
The private key and certificate for each of the Intranet servers are loaded into the SSL appliance because it is using known server key mode to decrypt the traffic. The [PT_OUTX2](#) diagram in [Passive-Tap Mode](#) on page 29 shows the connection mode being used in this example.

Inbound and Outbound Inspection

The following diagram shows a deployment where both inbound and outbound traffic are inspected. The IPS in this deployment can detect any threats in inbound sessions heading for the Intranet servers from users on the Internet and, at the

same time, can detect any inbound threats over sessions from users on the LAN to Internet servers.

Inbound and Outbound Inspection with IPS and Network Forensic Appliances




In addition, the Network Forensic system can detect and identify any files sent out as webmail attachments by internal users. In this example the SSL appliance uses both certificate re-sign and known server key mechanisms to decrypt traffic. The selection of which mode to use is determined by whether an SSL session is incoming or outbound. The [AI_FTA1xCP2T01](#) diagram in [Active-Inline Mode](#) on page 33 shows the connection mode being used in this example.

High Availability Deployment

An SSL appliance segment has fail-to-wire capabilities provided by the NetMod to ensure connectivity for most scenarios where hardware has failed or software is temporarily not available. However, some customers prefer to deploy multiple SSL appliances to ensure that in High Availability scenarios, traffic continues to be inspected. A typical High Availability deployment is depicted in the following diagram.

High Availability Deployment



The SSL appliance segment is configured in HA mode with the software failure mode set to **Disable Interfaces** and enabling link state mirroring on the ethernet switch devices. Normally switches A1 and A2, SSL appliance A, and its attached security appliances, are active. If any of the links along that path fail, or if the SSL appliance or its attached security appliance or either of the ethernet switches fail, the link down state will propagate with standard mechanisms like the Spanning Tree Protocol or the Virtual Router Redundancy Protocol ensuring that traffic is rerouted over the link between switches B1 and B2 that pass through SSL appliance B (dashed line in the figure). Availability can be further improved by including additional links between switches A1 and B1 and between switches A2

and B2 (shown as dashed lines in the diagram). This ensures that traffic can flow from Network X via A1 to B1 and then through SSL appliance B if required. Depending on the required availability levels and the built in redundancy features of the switches, devices A1 and B1 may be combined into a single device, with A2 and B2 being similarly combined.

Contact Sourcefire support should you require more information with respect to High Availability deployment options.

Chapter 3

Physical Installation

This section describes the following procedures:

- Installing the Sourcefire SSL appliance as a rack-mounted component.
- Connecting the Sourcefire SSL appliance to the network.

Safety Information

Follow the warnings and cautions listed in [Safety Information](#) on page 149 when installing or working with the SSL appliance.

WARNING! Read all the installation instructions before connecting the appliance to its power source. See the important safeguards in [Safety Information](#) on page 149 for information regarding the setup and placement of the SSL appliance.

Requirements Checklist

The following will be required:

Requirements Checklist

SSL Appliance 1500	SSL Appliance 2000	SSL Appliance 8200
At least 1U rack space (deep enough for a 27" device) – power and management ports at rear	At least 1U rack space (deep enough for a 27" device) – power and management ports at rear	At least 2U rack space (deep enough for a 27" device) – power and management ports at rear
Phillips (crosshead) screwdriver	Phillips (crosshead) screwdriver	Phillips (crosshead) screwdriver
Two available power outlets (110 VAC or 220-240 VAC)	Two available power outlets (110 VAC or 220-240 VAC)	Two available power outlets (110 VAC or 220-240 VAC)
Two IEC-320 power cords (that is, standard server or PC power cords) should the supplied power cords not be suitable for your environment	Two IEC-320 power cords (that is, standard server or PC power cords) should the supplied power cords not be suitable for your environment	Two IEC-320 power cords (that is, standard server or PC power cords) should the supplied power cords not be suitable for your environment
Cooling for an appliance with two 450W power supply units	Cooling for an appliance with two 650W power supply units	Cooling for an appliance with two 750W power supply units
One RJ-45 CAT5e/CAT6 Ethernet cable to connect the SSL Appliance 1500 to the management network (or a local notebook or desktop computer which is used to manage the SSL Appliance 1500)	One RJ-45 CAT5e/CAT6 Ethernet cable to connect the SSL Appliance 2000 to the management network (or a local notebook or desktop computer which is used to manage the SSL Appliance 2000)	One RJ-45 CAT5e/CAT6 Ethernet cable to connect the SSL Appliance 8200 to the management network (or a local notebook or desktop computer which is used to manage the SSL Appliance 8200)
Appropriate copper or fiber cables to connect the eight active interfaces to the network and to associated security appliances	Appropriate copper or fiber cables to connect NetMods to the network and to associated security appliances	Appropriate copper or fiber cables to connect NetMods to the network and to associated security appliances


Rack Mounting

The SSL appliance is equipped with installed rack mount brackets and supplied with rack mount rails allowing easy installation in a rack.

Rear

The rear of the SSL Appliance 1500 is shown below. Ventilation holes on the rear panel must not be blocked because free flow of air is essential for system cooling.

SSL Appliance 1500 Back Panel




The rear of the SSL Appliance 2000 is shown below. Ventilation holes on the rear panel must not be blocked because free flow of air is essential for system cooling.

SSL Appliance 2000 Back Panel



The rear of the SSL Appliance 8200 is shown below. Ventilation holes on the rear panel must not be blocked because free flow of air is essential for system cooling.

SSL Appliance 8200 Back Panel



SSL1500, SSL2000, and SSL8200 Back Panel Components

1	Serial Port	5	Management Ethernet 1
2	VGA Display Connector	6	Management Ethernet 2
3	USB Port	7	Power Supply 1
4	USB Port	8	Power Supply 2

The SSL appliance is equipped with two independent power supply units, either of which can power the appliance. The power supply units feature IEC-320 (that is, standard server / PC style) connectors. Attach both units to an uninterruptible power supply or other power outlet (110 or 220/240 Volt AC).

IMPORTANT! The power supplies are hot-swappable and can be replaced in while the SSL appliance is powered on and operating.

WARNING! Use only replacement units supplied by Sourcefire, Inc. Use of other units will void any warranty and may damage the system.

Front Panel

The SSL Appliance 1500 comes configured with eight copper or fiber interfaces. The SSL Appliance 2000 and SSL Appliance 8200 have modular I/O bays that allow for flexibility in the number of network interfaces and the type of media

supported. Network I/O Modules (NetMods) are installed in the bays to configure the desired combination of interfaces. 10Gig and GigE NetMods cannot be mixed in an SSL appliance chassis, so a device may either have GigE NetMods or 10Gig NetMods. There are three front-facing modular I/O bays in the SSL Appliance 2000. There are seven front-facing modular I/O bays in the SSL Appliance 8200. Available NetMod options are listed below; other NetMod types may become available in the future:

- 4 x GigE copper (4 ports of 10/100/1000Base-T with bypass)
- 4 x GigE fiber (4 ports of 10/100/1000Base-SX with bypass)
- 2 x 10Gig fiber (2 ports of 10GBase-SR with bypass)
- 2 x 10Gig fiber (2 ports of 10GBase-LR with bypass)

WARNING! The software supports a maximum of sixteen external interfaces. If 4 x GigE NetMods are used, a maximum of four can be installed in the system.

The image below is of an SSL Appliance 8200 with four NetMods installed, two 4 x GigE fiber interfaces, and two 4 x GigE copper interfaces.

SSL Appliance 8200 Front View




NetMods and the switch module installed in the front facing bays of the SSL appliance are NOT hot-swappable and are not user-replaceable items. **Do not** attempt to remove or insert NetMods or switch modules. Insertion or replacement of NetMods and switch modules should only be carried out by trained service personnel and should only be done when the system is powered off.


The front panel has indicators, buttons, an LCD display, and a USB port that the administrator can use to configure and diagnose the system. The relevant portion of the front panel is shown in the following illustration and the [Front Panel Components](#) table identifies the components. [Initial Configuration and Setup](#) on page 67 provides details on how the front panel components can be used to configure the system.

Note that the unit pictured in the following illustration has a 4 x GigE copper NetMod installed in the right-hand bay.


SSL Appliance 1500 Front Panel Components



SSL Appliance 2000 Front Panel Components



SSL Appliance 8200 Front Panel Components



Front Panel Components

1	Switch Module	7	System Status Indicator
2	Keypad Array	8	NMI Button (recessed)
3	LCD Display	9	Reset Button (recessed)
4	Management Ethernet 1 Indicator	10	Identify Button
5	Management Ethernet 2 Indicator	11	Power Button
6	Disk Activity Indicator	12	USB socket

The front panel status LEDs for the management Ethernet ports are green when the link is up and flash amber/yellow to indicate traffic flowing over the link. The two LEDs that are part of the Ethernet ports on the rear panel indicate the operating speed of the link and if data is flowing over the link. The left LED viewed from the back of the unit is green if the link is up and flashes to indicate traffic flow. The right LED can be: off, indicating a 10Mbps connection; green, indicating a 100Mbps connection; or amber, indicating a GigE connection.

The disk activity LED is green and flashes when there is any disk activity on a SATA port in the system.

The system status LED is green/amber and the various display options indicated different system states.

The [System Status Indicator Meaning for SSL1500](#) table shows the various system states that can be indicated by the system status LED on the front panel of the SSL Appliance 1500.

System Status Indicator Meaning for SSL1500

Color	State	System Status	Meaning
None	Off	OK	System ready - no errors detected
Red	Solid	Fault	AC power supply failure No AC power cord present Absence of AC power supply module

The [System Status Indicator Meaning for SSL2000 and SSL8200](#) table shows the various system states that can be indicated by the system status LED on the front panel of the unit.

System Status Indicator Meaning for SSL2000 and SSL8200

Color	State	System Status	Meaning
Green	Solid	OK	System ready - no errors detected
Green	Blink	Degraded	Memory, fan, power supply or PCIe failures
Amber	Solid	Fatal	Alarm – system has failed and shut down
Amber	Blink	Non-Fatal	Alarm – system likely to fail – voltage/temp warnings
Green + Amber	Solid	OK	First 30 seconds after AC power connected
None	Off	Power Off	AC or DC power is off

The NMI and Reset buttons are recessed, requiring the use of a straight thin object to press the button. Pressing the Reset button will cause the system to be reset. The NMI button should not be pressed during normal operation because it may cause the system to halt. If the NMI button is pressed, this fact will be recorded in the system log file.

If you press the ID button, a blue LED on the rear panel to the left of the serial port illuminates. This LED is located behind the back panel and is visible through the ventilation holes. This LED makes it easier to locate a system when it is racked with other systems.

Connecting to the Network

The SSL Appliance 1500 has eight front facing copper or fiber interfaces. The image below shows an SSL Appliance 1500 with eight copper interfaces.

SSL Appliance 1500 with copper interfaces



Ports are numbered from left to right when facing the front of the device. When a segment is configured and activated the port numbers allocated to that segment are displayed on the management WebUI. The relevant ports must be connected to the network and associated security appliance(s) using appropriate copper or fiber cabling.

The left LED at the top of the socket indicates link status and the right LED at the top of the socket indicates link activity. The left LED can be: off (indicating no connection), green (indicating a 1000Mbps connection), or amber (indicating a GigE connection).

Below each pair of interfaces is a Fail-To-Wire (FTW) status LED that indicates the current FTW status for that pair of interfaces.

Fail-To-Wire Status

Color	State	FTW Status
None	Off	Active State
Green	Solid	Active State with armed watchdog
Amber	Solid	Commanded FTW state change
Amber	Flashing	Forced FTW

The image below shows an SSL Appliance 1500 with eight fiber interfaces.

SSL Appliance 1500 with copper interfaces



Each fiber interface has two LEDs arranged vertically. The top LED indicates link activity and the bottom LED indicates link state. Link state can be off (indicating no link is established), or solid green (indicating a 1000Mbps link is established).

Each pair of fiber ports has a FTW LED that indicates FTW status as shown above.

IMPORTANT! Pairs of ports share fail-to-wire (FTW) hardware that is used to directly connect the two ports together whenever the port pair are in FTW mode. If the box is powered off then all ports will be in FTW mode so each pair of ports will be connected to each other.

The SSL Appliance 1500 comes configured with either eight copper or eight fiber interfaces.

The modular I/O bays in the SSL Appliance 2000 and SSL Appliance 8200 allow for flexibility in the number of network interfaces and the type of media supported. Network I/O Modules (NetMods) are installed in the bays to configure the desired combination of interfaces. 10Gig and GigE NetMods cannot be mixed in an SSL appliance chassis; a device may either have GigE NetMods or 10Gig NetMods. There are three front-facing modular I/O bays in the SSL Appliance 2000. There are seven front-facing modular I/O bays in the SSL Appliance 8200. Available NetMod options are listed below; other NetMod types may become available in the future:

- 4 x GigE copper (4 ports of 10/100/1000Base-T with bypass)
- 4 x GigE fiber (4 ports of 10/100/1000Base-SX with bypass)
- 2 x 10Gig fiber (2 ports of 10GBase-SR with bypass)
- 2 x 10Gig fiber (2 ports of 10GBase-LR with bypass)

[Known Server Key Decryption Method -Passive-Tap Mode](#) on page 21 shows an SSL appliance device with two NetMods installed, in this example the NetMods each support 2 x 10Gig fiber interfaces.

WARNING! NetMods are NOT hot-swappable and are not user-replaceable items. **Do not** attempt to remove or insert NetMods. Insertion or replacement of NetMods should only be carried out by trained service personnel and should only be done when the system is powered off.

Ports are numbered from left to right when facing the front of the SSL Appliance 2000, and left to right and top to bottom when facing the front of the SSL Appliance 8200. When a segment is configured and activated the port numbers allocated to that segment are displayed on the management WebUI. The relevant ports will need to be connected to the network and associated security appliances using appropriate copper or fiber cabling.

Chapter 4

Initial Configuration and Setup

The SSL appliance is configured and managed using a Web based User Interface (WebUI) which provides a graphical means to configure the device. The front panel keypad and display can be used to configure the management network settings for the device and are also used during initial bootstrap mode and to unlock the master key during system startup.

IMPORTANT! The SSL appliance is factory-configured to use DHCP to acquire an IP address for the management ethernet. The front panel keypad and LCD can be used to configure a different fixed IP address.

Bootstrap Phase

When the SSL appliance is powered on or re-booted, it goes through a number of stages before reaching the fully operational state. These stages are termed the *bootstrap* phase.

After the SSL appliance is powered on, it can be forced into one of three states by typing in the correct sequence on the front panel keypad. To enter factory default reset mode, the key sequence must be typed within 5 seconds of power on. Key sequences for other states can be typed at any time.

Enter code on keypad to enter one of three states:

- Factory default reset
- IP configuration mode
- PIN entry mode

The front panel keypad on the appliance has the keys arranged in the following layout:

0	1
2	3

The following key sequences are used to enter one of the three states described above.


Power On Sequences

Sequence	State Entered
031203	Factory default reset
01320132	IP configuration mode
01230123	PIN entry mode

Factory default reset and IP configuration mode can both be run before the system enters the main bootstrap phase. Pin entry mode can be run after the system enters the main bootstrap phase.

- Factory default reset causes the box to reset and erases all configuration and other data on the system, returning it to exactly the same state as when it was received from the factory.
- IP configuration mode allows the management network to be configured to use DHCP to get an address or to configure a static IP address. The IP address settings will then be used during the bootstrap phase and will be saved so it is used after the bootstrap phase is over.
- Pin entry mode is explained later in this section.

The following illustration shows the front panel LCD with the first screen of the menu to configure the management network settings. The four symbols at the right of the display indicate what each of the four buttons does.



The main sequence of events during bootstrap is shown below. Depending on the initial state of the SSL appliance, some of these steps may or may not apply:

- Choose Master Key Mode; this step only occurs if the mode is not already set.
- Find or create the master key; if the master key is password-protected, unlock using password.
- Create at least one user with the Manage Appliance role and one with the Manage PKI role if they do not exist. This step won't occur if there are already users with these roles.

All the above steps are managed using a limited version of the WebUI.

Configuring Static IP Address for Management

The easiest way to configure the SSL appliance is to allocate it a management IP address using DHCP. However, if a static IP address is required then it can be configured by interrupting the start up sequence using the keypad sequence **01320132** then using the front panel keypad and LCD to configure the desired address. On the initial screen you can enable or disable DHCP by pressing the top or bottom rightmost button on the front panel keypad.

To configure a static IP address use the up and down arrow buttons to move to screens that allow the address information to be configured. Press the down arrow key until you reach the **IP** option. Use the up/down arrow key to select the item to be configured and then press the top right button on the keypad to edit that item. The items that can be selected and configured are:

- IP address for the system
- IP Netmask for the system
- Gateway IP address for the system

After selecting an item to edit, the buttons are mapped to left and right arrow to allow the cursor to be moved within the item being configured and the up arrow key that is used to change the value at the point the cursor is located.

On entry to the IP configuration screen, the cursor is located under the leftmost digit in the address. The left/right arrow buttons will move the cursor. Press the up arrow button to increment the number above the cursor. After all the changes to the IP address are complete, press the top right button to exit back to the previous level in the menu. Configure other elements, such as the **Netmask**, in the same manner. Select **Apply** after you have configured all the elements.

Password Entry

The password used to unlock the master key must be typed in on the front panel keypad after entering the code for PIN entry mode. The password is required only if the master key mode chosen requires a PIN. The password is a minimum of eight characters long and the user must select each character from a set of four

characters that are displayed on the LCD. Passwords can include upper and lower case characters and the space character. The mechanism used to enter a password is described below.

Characters are selected using the buttons on the keypad and four button presses are required to input each character in the password. Each button press narrows down the set of characters that can be selected with the final button press choosing a specific character.

The first menu option allows for selection of upper or lower case for the character being entered. The three remaining menus narrow down the selection of the character to be input. The second menu allows for selection of a character group with the letters **A**, **J**, or **S** identifying the character group as shown on the grid below.

A	D	G	J	M	P	S	V	Y
B	E	H	K	N	Q	T	W	Z
C	F	I	L	O	R	U	X	-

Choosing a character limits future selection options to other characters that are the same color in the grid. The third menu allows the selection of a subset of the character group already selected with the subset being identified by either **ADG** or **JMP** or **SVY** depending on which character was selected from menu 2.


The final menu allows selection of the character to be used in the password from the three characters in the vertical column with the character selected from menu 3 at the top. So, if **A** was chosen from menu 3 then menu 4 will offer the characters **A**, **B**, and **C**.

Note that the bottom character in the column with **Y** at the top is the space character.

The following sequence of images shows the LCD display at various points during the process of entering the password **Pass word**.


The following graphic shows the initial menu display after PIN entry mode is active. The four characters at the right of the display correspond to the four buttons with the two upper buttons being used to select upper or lower case for the character. The lower left button is a backspace key to erase a selection and the lower right button is used to enter the chosen selection.

PIN Entry Menu 1 -Select Upper or Lower Case




The next graphic shows the second menu in the PIN entry process, which allows selection of the group of characters that will be used. Notice that the characters are shown in upper case because this was selected on the preceding menu. Because the password in this example is **Pass word**, select **J** from the grid shown earlier. **P** is part of the green block of characters which includes **J** at the top left of the block.

PIN Entry Menu 2 -Select the Character Group




The next graphic shows the third menu in the PIN entry process which allows selection of the sub group of characters to be used. In this example the character we want is **P** and this is shown as an option. Note however that selecting **P** in this menu indicates the sub group containing the characters **P, Q, and R**.

PIN Entry Menu 3 -Select the Character Subgroup



The next graphic shows the fourth and final menu in the PIN entry process which allows the desired character to be selected. In this example, the character **P** is selected by pushing the top left button in the keypad.


PIN Entry Menu 4 -Select the Character



The next graphic shows the display after the first character in the password has been entered. Note that the system is now back at menu 1 in the process


allowing the choice of upper or lower case to be selected for the next character in the password.

PIN Entry -First Character Entered




The next three graphics show the steps in the process of entering the second character in the password.


PIN Entry -Menu 2 -Select the Character Group



PIN Entry -Menu 3- Select the Character Subgroup




PIN Entry -Menu 4 -Select the Character




The next graphic shows how the space character is entered into the password. Use the bottom left button to select the space character, shown as a space on the LCD display.

PIN Entry -Menu 4 -Select the Space Character




The next graphic shows the space character in the partially entered password.

PIN Entry -Menu 4 -Showing the Space Character in the Display



The next graphic shows the final complete password which is saved by pressing the bottom right button. After the password has been entered and accepted, it is stored in the system and is used at the appropriate point in the bootstrap sequence.

PIN Entry -Menu 4 -Showing the Completed Password



Installation Process

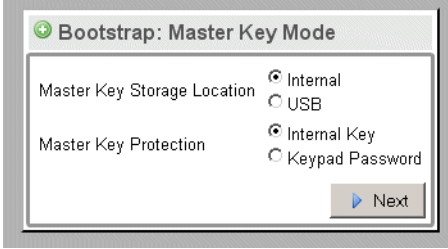
A typical installation process for a new SSL appliance is as follows:

- Install the system in a rack in the equipment room.
- Power the appliance up and use the keypad to enter IP configuration mode, and configure a valid address for the device.
- Use the keypad to enter PIN entry mode and enter a PIN. Make a note of the PIN.
- Use a system that allows you to access the WebUI and complete the WebUI part of the bootstrap process. Note that a new SSL appliance will not have an existing master key and one will be created. After the master key has been created, you must re-enter the PIN via the front panel keypad.

The first step (as shown in the following graphic) occurs only if the master key mode is not already configured. If the master key mode is configured, this step will not occur. This allows configuration of the location where the Master Key for the SSL appliance is stored and whether or not it is password protected. For the highest level of security, part of the Master Key can be stored on an external USB memory device and can be password protected. You will need the USB memory

device when the device is powered on and you must input the password on the front panel keypad to make the device operational.

Bootstrap Master Key Mode

A screenshot of a web-based configuration window titled "Bootstrap: Master Key Mode". The window has a light gray border and a title bar with a green plus icon. Inside, there are two sections: "Master Key Storage Location" and "Master Key Protection". Under "Master Key Storage Location", there are two radio buttons: "Internal" (selected) and "USB". Under "Master Key Protection", there are two radio buttons: "Internal Key" (selected) and "Keypad Password". At the bottom right of the window is a button labeled "Next" with a blue play icon.

After the master key mode is configured, the appliance will scan the internal and, if required, external persistent storage device for the master key. If the master key is not found the system will create the master key. If the master key is protected by a password, the user must first enter the password on the keypad before the master key can be unlocked or created. While in this state, the WebUI will display a screen with a *spinner* and without any buttons or links.

IMPORTANT! The password can be entered into the device before the WebUI bootstrap phase in which case it will be retrieved and used when this point in the bootstrap sequence is reached.

After the master key is unlocked, the secure store can be opened or created.

The final stage of the bootstrap process is user setup. At least one user with the Manage Appliance role and at least one user with the Manage PKI role must be created. You can create one user with both roles, or create separate users for these roles. After the users are created, the WebUI will go to the login screen, after which the user can log in with real credentials and configure the SSL

appliance. The screen allowing configuration of users with these roles is shown in the following graphic.

Bootstrap User Setup

Bootstrap: User Setup

Manage Appliance user required. Manage PKI user required.

User ID:

Full Name:

Roles:

Password:

Confirm Password:

IMPORTANT! If the system has previously been configured and has at least one user with the Manage Appliance role and one user with the Manage PKI role, this step will be skipped.

After creating the necessary users the normal system login screen will appear, allowing the user to login, at which point they will have access to the full WebUI to manage the SSL appliance. A user with the Manage Appliances role can now create additional users but cannot give these users the Manage PKI role. Only a user with the Manage PKI role can give this role to a user.

Whenever the SSL appliance is powered on or forced into a factory default reset, the bootstrap phase will run before the device becomes fully functional. Depending on how the device is configured, the administrator may need to provide input to enable the bootstrap phase to complete before the device is operational again.

- If the master key is stored internally and no password is set for the master key, the bootstrap process becomes invisible and the device will start up without any need for input from the administrator.
- If the master key is partly stored on a USB storage device, this storage device must be connected to the system before the bootstrap phase can complete.
- If the master key is protected by a password, you must enter the password using the front panel keypad before the bootstrap phase can complete.
- If the master key is partly stored on a USB storage device and is protected by a password, you must enter the password using the front panel keypad and the USB storage device must be connected before the bootstrap phase can complete.

Network Connections

HTTPS and SSH access to the SSL appliance is via the separate management Ethernet interface and should be connected to a secure network used by administrators to manage security appliances. Connect Management Ethernet 1 to the secure management network (see the [SSL1500, SSL2000, and SSL8200 Back Panel Components table](#) on page 59).

By default, the SSL appliance uses DHCP to acquire an IP address from the network. The acquired address can be viewed on the front panel LCD. If DHCP is not in use, a static IP address can be configured. See [Bootstrap Phase](#) on page 67.

The front panel keypad and LCD display can be used to change the management network settings before or after the bootstrap process. Use the menus on the LCD to configure the SSL appliance to use DHCP to obtain an IP address or to configure a fixed IP address.

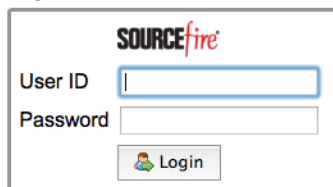
Post Bootstrap Configuration

After the bootstrap phase is complete, the full WebUI is available and can be used to configure the system. The WebUI is described in detail in [Web-Based Management Interface \(WebUI\)](#) on page 102. This section provides a quick summary of the basic configuration steps. An HTTPS connection to the IP address assigned to the SSL appliance management interface produces the standard login box.

IMPORTANT! The SSL appliance uses a self-signed SSL server certificate which may result in a warning message from the browser when connecting to the WebUI. The warning can be prevented by adding this self-signed certificate to your browser as a trusted device. Consult your browser documentation for details on how to add the SSL appliance as a trusted device.


The following graphic shows the login box which appears in the center of the initial access screen.

Login Box on the Initial Access Screen

A login box with the Sourcefire logo at the top. It contains two input fields: 'User ID' and 'Password'. Below the password field is a 'Login' button with a user icon.

The bottom of the initial access screen displays additional information on the SSL appliance.

Status Information on the Initial Access Screen



This status information allows you to determine what version of software the SSL appliance is running without the need to log into the system. The following graphic shows the top and bottom of the initial management dashboard screen after the administrator has logged on. The top of the screen contains menus on both the left and right side. The names of the two menus on the right side depend on the device name and the username.

Initial Management Screen After Login



In this example, the SSL appliance has a device name of `localhost.localdomain` and the username of the connected user is `admin`. The bottom of the screen contains status information on the device and shows:

- current date and time
- version of software running on the device
- status indicators for System, Load and Network

The status indicators change color if there are problems. Alert messages occur in a panel above the bottom status line. None are present in the example screen.

As part of an initial configuration, the following would normally be configured:

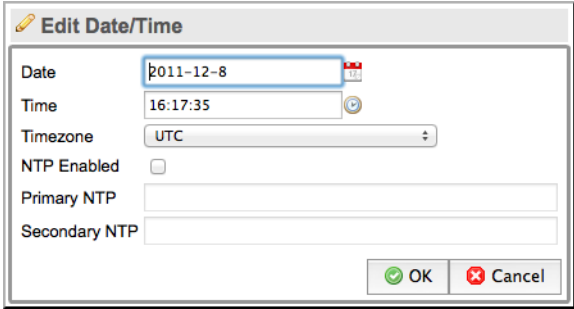
- Management network settings
- Time zone and use of NTP
- Additional user accounts with relevant roles assigned to the user

Configuring System Date/Time and Time Zone

To configure the system date and time, use the **Date/Time** option on the Device menu. In the previous graphic, the Device menu is labeled **localhost.localdomain** because that is the example systems name.

Click the pencil icon at the top right of the Date/Time box to edit these settings. The following graphic shows the edit screen and settings that can be changed.

Date and Time Configuration



Edit Date/Time

Date: 2011-12-8

Time: 16:17:35

Timezone: UTC

NTP Enabled: ☐

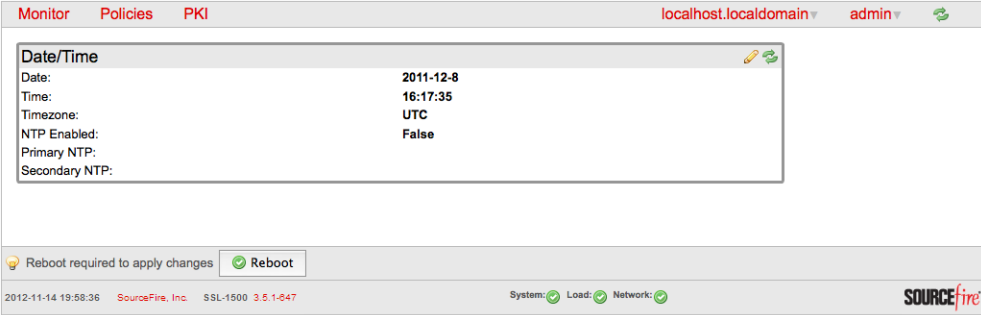
Primary NTP:

Secondary NTP:

OK Cancel

If NTP is enabled, as in this example, the Date and Time input boxes will be disabled. These values are set by the Network Time Protocol (NTP). To use NTP to operate, you must configure a primary NTP server and, ideally, a secondary NTP server. After the settings are configured, click **OK** to save the settings. The following graphic shows the updated Date/Time box.

Time Settings with the Reboot Button



Monitor Policies PKI localhost.localdomain admin

Date/Time

Date: 2011-12-8

Time: 16:17:35

Timezone: UTC

NTP Enabled: False

Primary NTP:

Secondary NTP:

Reboot required to apply changes Reboot

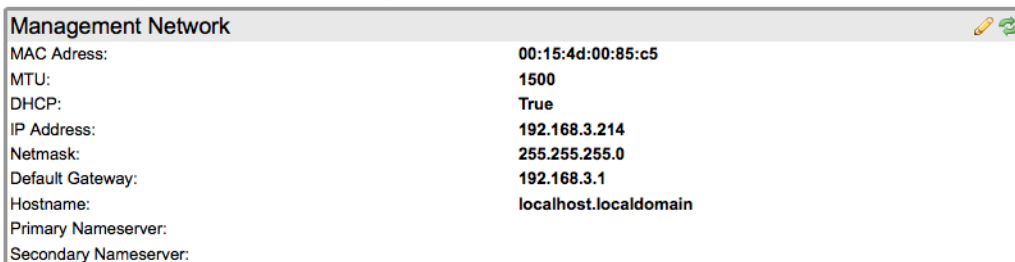
2012-11-14 19:58:36 SourceFire, Inc. SSL-1500 3.5.1-047 System: Load: Network: SOURCEfire

To change the time, click **Reboot**. This reboots the system.

Configuring Management Network Settings

Configure the management network settings using the Management Network menu option on the device menu. The following graphic shows the details displayed by this option.

Management Network Settings

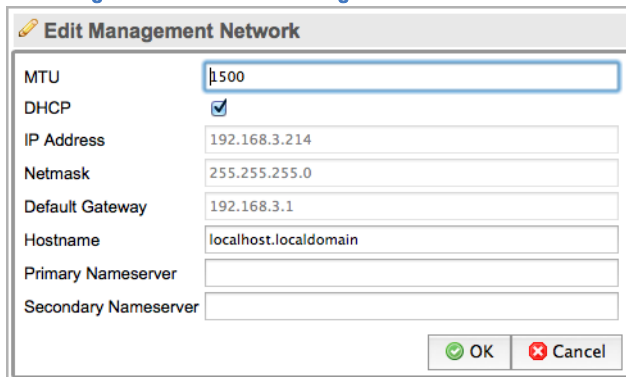


The Management Network Settings window displays the following configuration details:

MAC Address:	00:15:4d:00:85:c5
MTU:	1500
DHCP:	True
IP Address:	192.168.3.214
Netmask:	255.255.255.0
Default Gateway:	192.168.3.1
Hostname:	localhost.localdomain
Primary Nameserver:	
Secondary Nameserver:	

Click the pencil icon at the top right to edit these settings. The following graphic shows the configuration screen and the parameters that can be edited.

Edit Management Network Settings Box



The Edit Management Network Settings Box displays the following configuration details:

MTU	<input type="text" value="1500"/>
DHCP	<input checked="" type="checkbox"/>
IP Address	<input type="text" value="192.168.3.214"/>
Netmask	<input type="text" value="255.255.255.0"/>
Default Gateway	<input type="text" value="192.168.3.1"/>
Hostname	<input type="text" value="localhost.localdomain"/>
Primary Nameserver	<input type="text"/>
Secondary Nameserver	<input type="text"/>

Buttons:

In this example, the system is configured to use DHCP to obtain an address. The IP Address, Netmask and Default Gateway input boxes are disabled. If DHCP is disabled, these fields will be editable.

This screen also allows configuration of SNMP parameters and the option to enable or disable SNMP management. The SSL appliance supports the standard SNMP MIB2 tables and uses the SNMP v2c version of the protocol. In order to allow SNMP management of the SSL appliance, enable SNMP and configure the SNMP parameters appropriately for your SNMP management system.

Press **OK** to save the settings. The screen will appear as shown below.

Network Management Network Settings with the Apply Button

Clicking **Apply** will cause a **Reboot** button to display. Changes to the network settings take place after the reboot has occurred.

Configuring Management Users

You can create additional user accounts on the system using the Users option on the platform menu. Click on the plus icon to add a new user to the system. The following graphic shows the Add User input box and the details required to add a user.

Add User Input Box

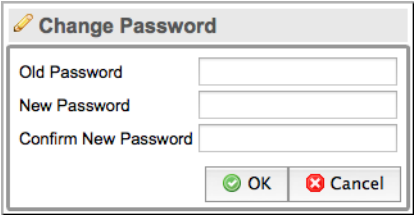
The Roles section of the input box allows assignment of one or more roles to the user being created. To assign more than one role, Ctrl-click the additional roles and then click **Save**. After you click **OK**, the new user is created and added to the system. The following graphic is displayed:

Current Users Configured in the System Display

User Management		
User ID	Full Name	Roles
admin	admin	Manage Policy, Manage Appliance, Auditor, Manage PKI

A user can change their own password at any time by logging on to the system and using the Change Password option on the User menu. The user menu is at the top right of the screen and displays the user's name as its title. A dialog box, as shown in the following graphic, allows the user to change their own password.

User Password Change Box

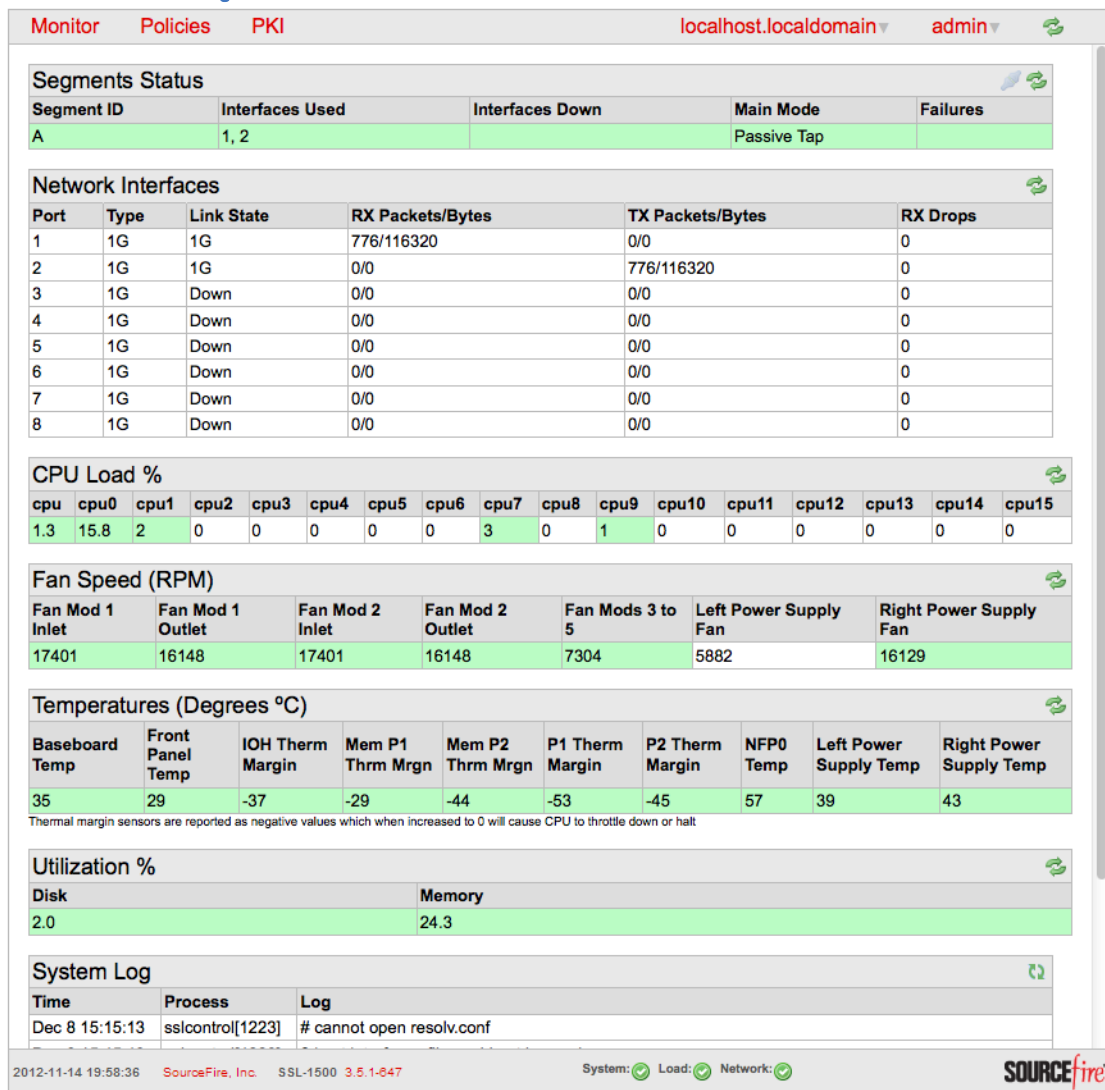
A screenshot of a 'Change Password' dialog box. The dialog has a title bar with a pencil icon and the text 'Change Password'. Inside, there are three text input fields labeled 'Old Password', 'New Password', and 'Confirm New Password'. At the bottom right, there are two buttons: 'OK' with a green checkmark icon and 'Cancel' with a red X icon.

Change Password	
Old Password	<input type="text"/>
New Password	<input type="text"/>
Confirm New Password	<input type="text"/>
<input type="button" value="OK"/> <input type="button" value="Cancel"/>	

System Status

View the overall status of the appliance by clicking on the [Monitor/Dashboard](#) menu option. The following graphic shows an example of the dashboard screen providing detail on the system status.

Management Dashboard Screen




Status details shown here feed into the summary status indicators for System, Load and Network that appear in the bottom line of the display.



Installing a CA for Certificate Re-Sign

Before the SSL appliance can be used to inspect traffic using Certificate Re-sign mechanisms the SSL appliance must have at least one CA certificate and private key installed for the re-signing. A CA can either be created by the SSL appliance (and self-signed or sent off for signing by another CA) or can be imported. If the SSL appliance has more than one CA for re-sign installed, it is possible to use different CAs to re-sign different SSL sessions by choosing the appropriate CA in the policy configuration. Management of Internal Certificate Authorities is done using the menu option on the PKI menu. The following graphic shows the screen when there are no Internal Certificate Authorities in the system.

Internal Certificate Authority Screen with No Entries



The icons at the top right allow the user to:

- Generate a new Internal Certificate Authority 
- Add an Internal Certificate Authority by importing an existing CA and key 

The subsections that follow consider each of these ways of adding an Internal Certificate Authority.

Creating a CA

Click on [Generate CA](#) icon to display the input form shown in the following graphic.

Generate Internal Certificate Authority Input Box

The screenshot shows a "Generate Certificate" dialog box. It contains several input fields: "Common Name" (demo-Sourcefire), "Division/Department/Org. Unit" (IT), "Company/Organization" (Sourcefire), "City/Town/Locality" (Columbia), "Country Code" (United States), "State" (MD), "Key Size" (1024-bit), and "Valid For" (5 years). At the bottom, there are three buttons: "Generate self-signed CA", "Generate certificate signing request", and "Cancel".

This allows the basic data required in a CA to be input and the key size and validity period to be specified. After the data is input, you have two options:

- Generate a self-signed CA
- Generate a certificate signing request (CSR)

If you choose the option to generate a self-signed CA, no further steps are required; the CA is generated and added to the set of Internal Certificate Authorities in the system. Because this CA is self-signed, it will not be trusted by client systems until it has been exported and added to the list of trusted CAs on the client system. See [PKI Management](#) on page 126 for details on how to do this. When the OK button is clicked the certificate is saved and installed and an entry in the Internal Certificate Authorities table appears with an indication that no CSR has been generated for this certificate.

If you choose the option to generate a CSR, a PEM format CSR is generated and must be sent to the Certificate Authority that is going to sign it. The following graphic shows an example CSR.

Internal Certificate Authority



1 Certificate Signing Request

Certificate Signing Request PEM:

```
-----BEGIN CERTIFICATE REQUEST-----
MIIBqTCCARICAQAwTEYMBYGA1UEAxMPZGVtby1Tb3VyY2VmaXJlMRMwEQYDVQK
EwpTb3VyY2VmaXJlMQswCQYDVQQLAwJVVDERMA8GA1UEBxMIQ29sdWliaWExCzAJ
BgNVBAGTAK1EMQswCQYDVQQGEwJVUzCBnzANBgkqhkiG9w0BAQEFAAOBjQAwYkC
gYEAvez/5a+CzWxsqMz1lAzXpQCPKG+Styydtia3OYti1jgq6Sj/GBh1Ngct050is
Wpd8hl1fUKMSajLZG7zkHuxvu4taDrgMzsZVYJHegNqXlM/PAGJI7KCCnZtNGyhnX
Sf5untNhoOmI/NF7MIRzzJLceffBqMVLACHNUUGx6ZtOaq0CAwEAaAAMA0GCSqG
SIb3DQEBBQUAA4GBAISRczB4JmAwMB0IfAeCy9dMdYFFK18kXWZgxR5wV6+HHpHN
rGVxj8IMb0ul48DsZY8q1WaW7uYbf3t1c1aJs263AdDmhQPmh3xoMKqtshxzJ00
inNDaHHQwk0gRtPe66Ju71PKk9PlwJDzbqGOy+tzboHgEife12VvNlthi0o2
-----END CERTIFICATE REQUEST-----
```

OK

The text in the CSR box should be copied into a file and that file then must be communicated to the CA that will sign the final Internal Certificate Authority certificate. When the OK button is clicked, the certificate details are saved and an entry in the Internal Certificate Authorities table appears with an indication that a CSR has been generated for this certificate. At this point, the certificate is not installed in the system because the signed Internal CA has not been received back from the CA that is signing it. When an entry in the table shows CSR True,

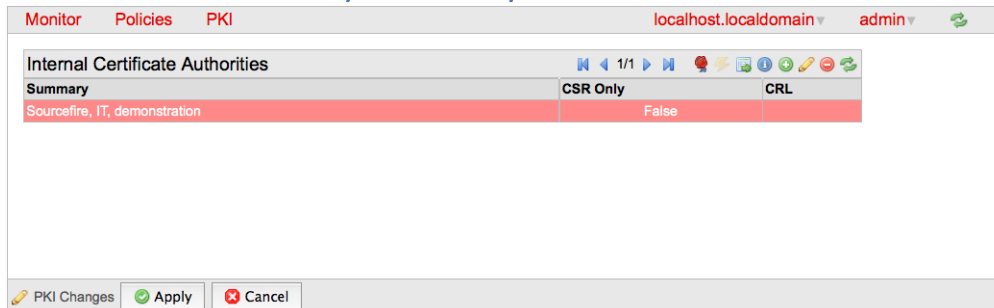
the icon to install a certificate is active and, if used, will prompt the user to provide the signed CA so it can be installed in the system.

WARNING! It is important to understand that the CSR is for a Certificate Authority and not for a normal SSL server certificate. The CA that will be used to sign this certificate will, in almost all cases, be the root CA of a private PKI domain and **not** a public CA. If the organization has a private PKI domain and client machines in the organization are configured to trust the private root CA, the CSR must be presented to the private root CA. The private root CA must sign the CSR to create a private Intermediate CA. The private Intermediate CA must then be loaded on the SSL appliance. Because the private Intermediate CA is signed by a root CA they trust, client machines will trust the private Intermediate CA.

WARNING! Public Certificate Authorities will sign CA CSR requests to create Intermediate CAs that are publicly trusted but there are onerous conditions and significant costs involved in doing this.

After the CSR has been generated, the Internal Certificate Authority screen will look like the following graphic:

Internal Certificate Authority with CSR Entry



Internal Certificate Authorities		
Summary	CSR Only	CRL
Sourcefire, IT, demonstration	False	False

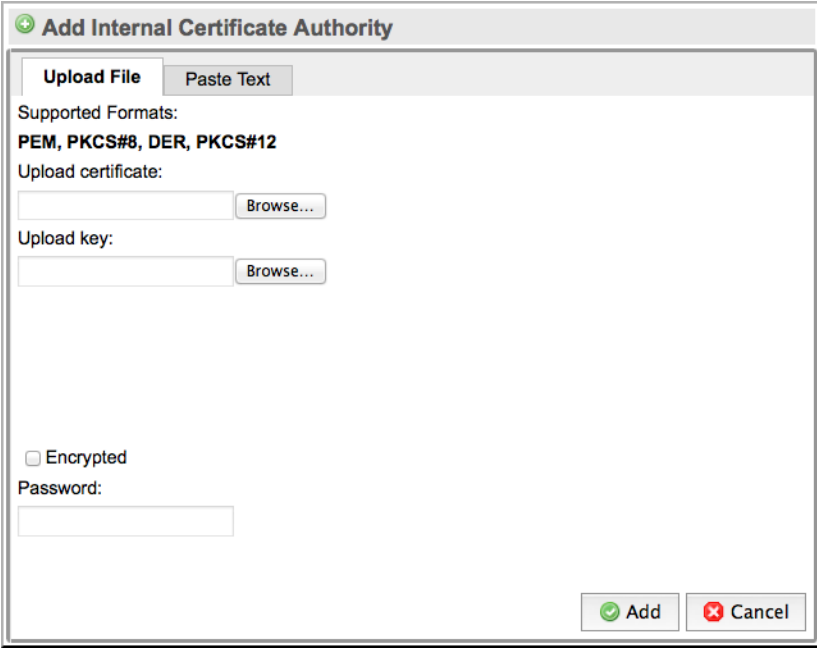
At this point, the CA cannot be used because the signed certificate from the CA that the CSR was sent to has not been loaded. After the signed certificate is available, it can be loaded by selecting the entry in the Internal Certificate

Authority box and then clicking on the ⚡ icon. This produces a box similar to the graphic in the next section, which allows the signed certificate to be imported into the system.

Importing a CA

If you already have a CA that you want to use as an Internal Certificate Authority in the SSL appliance, you can import this and install it in the system. You will need both the CA certificate and the private key for the CA to install it on the system. Click [Add](#) to generate a form that allows you to either select the files containing the certificate and private key or to paste in the certificate and private key directly. The following graphic shows the form used to import a CA.

Internal Certificate Authority Import Box



The screenshot shows a web form titled "Add Internal Certificate Authority". It has two tabs: "Upload File" (selected) and "Paste Text". Below the tabs, it lists "Supported Formats: PEM, PKCS#8, DER, PKCS#12". There are two sections for uploads: "Upload certificate:" with a text input and a "Browse..." button, and "Upload key:" with a text input and a "Browse..." button. At the bottom, there is a checkbox labeled "Encrypted" and a "Password:" text input. In the bottom right corner, there are two buttons: "Add" (with a green checkmark icon) and "Cancel" (with a red X icon).

If the certificate and key being imported have been encrypted and protected with a password, you will must check the [Encrypted](#) box and then type in the password in the [Password](#) box.

Importing Known Server Keys

To inspect traffic to an internal SSL server, the easiest approach is to use a known server mode which requires that a copy of the server's SSL certificate and private key, or just the private key, are loaded into the SSL appliance. Known server certificates and keys are imported into the [all-known-certificates-with-keys](#) list or the [all-known-keys](#) list and can then be copied to custom lists, if required. Use the [Known Certificates and Keys](#) option on the [PKI](#) menu to import new certificates and keys. The [Known Keys List](#) option on the [PKI](#) menu is used to import new keys.

There are two input forms provided: one chooses the list that is to be operated on, and the other manipulates the contents of that list. Initially, there is only one

list called `all-known-certificates-with-keys`, and it will not contain any certificates. The following graphic shows the initial appearance of the input forms:

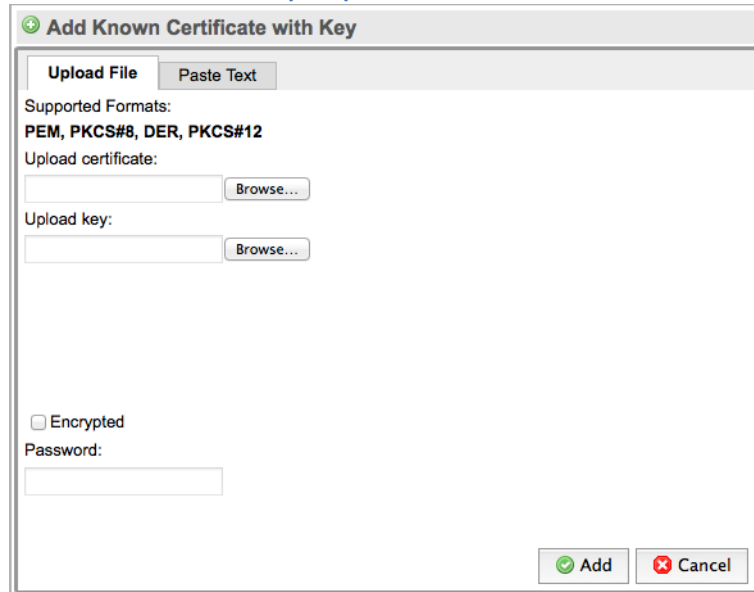
Known Certificate with Keys Display



The screenshot shows a window titled "Known Certificates with Keys Lists". It contains a table with one row: "all-known-certificates-with-keys". Below the table is a "Summary" section.

To import the first known server key and certificate, click on the [all-known-certificates-with-keys](#) entry under Known Certificates with Keys List and then click **Add** in the Known Certificate with Keys form. The following graphic shows the input form that will appear.

Known Certificate with Keys Import Box




The screenshot shows a dialog box titled "Add Known Certificate with Key". It has two tabs: "Upload File" and "Paste Text". Under "Supported Formats:", it lists "PEM, PKCS#8, DER, PKCS#12". There are two sections: "Upload certificate:" with a text field and a "Browse..." button, and "Upload key:" with a text field and a "Browse..." button. Below these is a checkbox labeled "Encrypted" and a "Password:" field. At the bottom right are "Add" and "Cancel" buttons.

You can either specify the files to import, or paste in the key and certificate details and click **Add**. If the key and certificate are valid, a message confirming that the Certificate has been added appears with a button that allows you to view the details of the imported certificate. The key now appears as a row in the Known Certificate with Keys form. A maximum of 8192 known server key/cert pairs can be loaded into the system.

The following graphic shows the screen after a number of keys have been imported. You must click the **Apply** button to save the imported certificates and keys to the secure store.

Known Certificate and Keys Display with Entries



[PKI Management](#) on page 126 explains how to create custom lists of Certificates and Keys in more detail.

Example Passive-Tap Mode Inspection


The following example shows the steps needed to configure the SSL appliance to inspect traffic that is intended for a server where you can obtain a copy of the private key and certificate. In this example the SSL appliance is deployed in passive-tap mode as described in [Passive-Tap Mode](#) on page 29. The known server certificates and keys used in this example are those shown in the graphic in the previous section.

The steps involved are:

1. Load the server key/certificate into the SSL appliance (see [Importing Known Server Keys](#) on page 86).
2. Create a ruleset that contains a rule to inspect traffic to the server.
3. Create a segment for passive-tap operation.
4. Activate the segment to start inspection.

In this example, the certificate and key for `test-web-server` is used to allow inspection of traffic going to that server. Because this certificate/key is already loaded into the system, you can proceed to the next step which is to create a ruleset that contains a rule specifying that traffic to `test-web-server` should be inspected. This is a two-step process, first creating the ruleset to hold the rule and then defining the rule itself. The following graphic shows the screen while adding a new ruleset called `passive-tap-example`.

Adding a Ruleset



After clicking **OK**, the new entry will appear as a row in the Rulesets grid and is available for use. A Policy Changes notification block with buttons to apply or

cancel the change appears at the bottom of the screen. Click **Apply** to complete the process and to save the ruleset to disk.

Click on the **passive-tap-example** row to select it. This displays Ruleset Options for this ruleset. In this example, the default settings need no changes and are explained below:

- No Internal Certificate Authority because we are not re-signing certificates
- All External Certificate Authorities and CRLs are used when checking an SSL session
- No trusted certificate are being used for systems that either have self-signed certificates or certificates signed by untrusted Certificate Authorities. If there were trusted certificates loaded into the system, the default setting would be to use All Trusted Certificates.
- Any SSL sessions that do not match a rule in this ruleset will be cut through to the attached security appliance without being decrypted

Click **Add** in the Rules grid section to display the Insert Rule form. Selecting **Decrypt (Certificate and Key known)** from the drop-down menu in this form will allow the valid options to be configured for this rule. The following graphic shows this form with the data entered.

Add Rule to Decrypt Using Known Server Key/Certificate

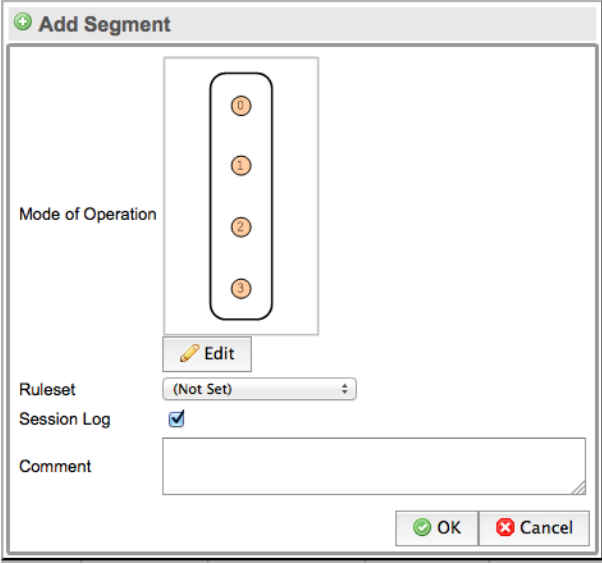
The screenshot shows the 'Insert Rule' dialog box. The 'Action' dropdown is set to 'Decrypt (Certificate and Key known)'. The 'Comment' field is empty. Under 'Known Certificate with Key', the dropdown is set to 'Sourcefire, IT, test-web-server'. The 'Source IP' and 'Destination IP' options are selected with radio buttons. The 'Source IP List' and 'Destination IP List' options are set to '(Not Set)'. The 'Destination Port' field is empty. At the bottom right are 'OK' and 'Cancel' buttons.

In this example, the rule applies only to a single server for which the certificate and key are known. The **Known Certificate with Key** option is checked and the system with the loaded key is selected from the drop-down menu. You can add a comment to the Comment box, but no other options are used in this rule. Click **Save** to create the rule. At the bottom of the screen is the Policy Changes notification block with buttons to apply or cancel the change. Click **Apply** to complete the process and to save the rule to disk.

The final part of the process is to create a segment, configure the segment to use the ruleset just created, and then to activate the segment. To create a segment, go to the **Policies / Segments** menu option for the Segments information. Initially

there are no segments configured in the system. To create a new segment, click **Add** in the Segments table. The following graphic shows the initial form.

Add Segment Box

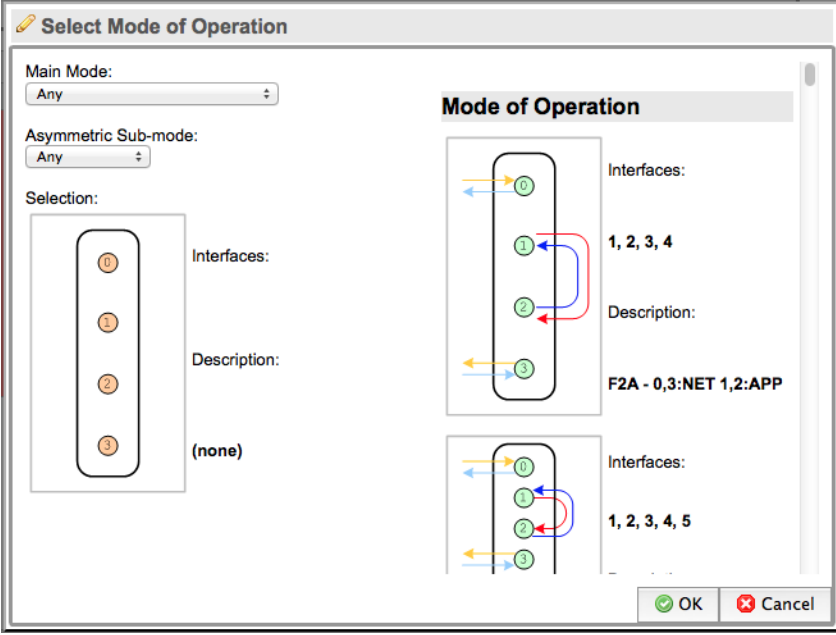


The 'Add Segment' dialog box contains the following fields:

- Mode of Operation:** A vertical list of four orange circles numbered 0, 1, 2, and 3.
- Edit:** A button with a pencil icon.
- Ruleset:** A drop-down menu showing '(Not Set)'.
- Session Log:** A checked checkbox.
- Comment:** A text area.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom right.

Click **Edit** to select the Mode of Operation for the required mode and then choose from the **Select Mode of Operation**. Choose the Ruleset from the drop-down menu. The following graphic shows the form used to select the mode of operation for a segment.

Selecting Mode of Operation for a Segment

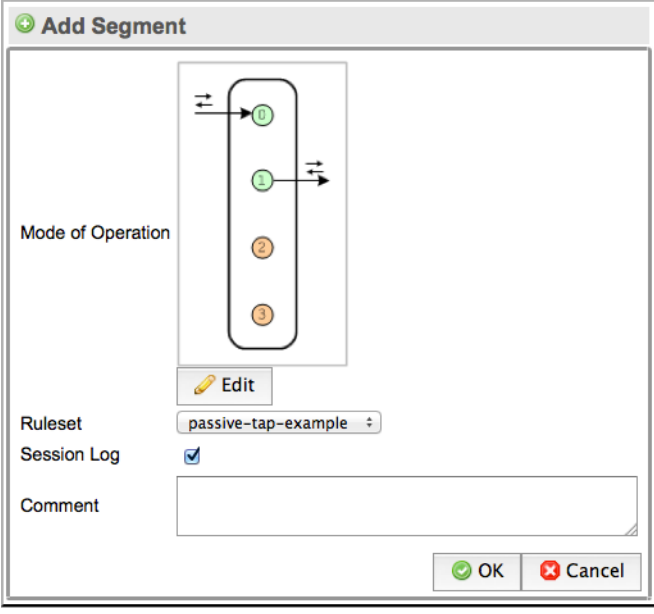


The 'Select Mode of Operation' dialog box contains the following fields:

- Main Mode:** A drop-down menu showing 'Any'.
- Asymmetric Sub-mode:** A drop-down menu showing 'Any'.
- Selection:** A vertical list of four orange circles numbered 0, 1, 2, and 3.
- Interfaces:** A text field showing '(none)'.
- Description:** A text field showing '(none)'.
- Mode of Operation:** A scrollable list of two diagrams. Each diagram shows a vertical list of four green circles numbered 0, 1, 2, and 3. The first diagram has arrows indicating traffic flow between interfaces 0, 1, 2, and 3. The second diagram has arrows indicating traffic flow between interfaces 0, 1, 2, 3, and 4.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom right.

The Mode of Operation part of the form has a scroll bar and displays all the different operating modes as images. The **Main Mode** drop-down menu allows the set of operating modes to be narrowed by choosing only passive-tap for example. This will reduce the number of options displayed in the Mode of Operations part of the form. The **Asymmetric Sub-Mode** drop-down menu can be used to further narrow the number of modes of operation that are displayed. Click on the image for the desired operating mode to select it and click **Save** to set this as the mode of operation for the segment. The following graphic shows the completed segment details before they are saved.

Passive-Tap Example Segment Configuration



The image shows a dialog box titled "Add Segment". Inside, there is a "Mode of Operation" section with a vertical list of four icons. The first icon is green and labeled "1", the second is green and labeled "2", the third is orange and labeled "3", and the fourth is orange and labeled "4". The first icon has a double-headed arrow pointing to it from the left. The second icon has a double-headed arrow pointing to it from the right. Below the icons is an "Edit" button. Below the "Edit" button is a "Ruleset" dropdown menu showing "passive-tap-example". Below the "Ruleset" dropdown is a "Session Log" checkbox which is checked. Below the "Session Log" checkbox is a "Comment" text box. At the bottom right of the dialog box are "OK" and "Cancel" buttons.

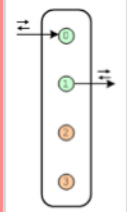
In this example, the session log has been enabled and the segment is using the **passive-tap-example** ruleset created earlier in the process. The graphic in the input box indicates that this segment will use two ports on the system. The actual ports used are determined when the segment is activated.

Click **OK** to create the segment. A Policy Changes notification block with buttons to apply or cancel the change will appear at the bottom of the screen. Click **Apply** to complete the process and to save the rule to disk.

The created segment is displayed in the Segments table and, as shown in the following graphic. Click on the segment to select it.

Passive-Tap Segment Options and Activations

System Options	
Overload Action:	Cut Through

Segments					
Mode of Operation	Segment ID	Ruleset	Interfaces	Session Log	Comment
	A	passive-ruleset	1, 2	Enabled	

Undecryptable Actions	
Compression:	Cut Through
SSL2:	Cut Through
TLS 1.2:	Cut Through
Diffie-Hellman in Passive-Tap mode:	Cut Through
Client Certificate:	Reject
Cipher Suite (including Export):	Cut Through
Uncached:	Cut Through

Certificate Status Actions	
Invalid Issuer:	(Not Set)
Invalid Signature:	(Not Set)
Expired:	(Not Set)
Not Valid Yet:	(Not Set)
Self Signed:	(Not Set)
Revoked:	(Not Set)
Status Override Order:	Rule over Segment

Plaintext Marker	
Type:	(Not Set)

Failure Mode Options	
Software Failure Action:	Fail-to-wire (Auto Recovery)
High Availability:	Disabled

There are three panels below the Segment panel in this table, each of which allows different types of actions to be configured for the selected segment. To change the settings in the Undecryptable Actions, Certificate Status Actions or Plaintext Marker panels, click [Edit](#) on that panel.

The Undecryptable Actions panel controls what happens to an SSL session that cannot be decrypted by the SSL appliance. Different actions can be configured depending on why decryption is not possible. In the example shown in the previous graphic, the action is to cut through the session except in the case where client certificates are used when the SSL session will be rejected.

The Certificate Status Actions panel-controls what happens if the server certificate used by the SSL session has particular errors in it. In this example, the action is to cut through the session for all error conditions. The Status Override Order line allows configuration of which Certificate Status actions have priority (those configured for the segment or those configured in a rule in the ruleset being used by this segment). In the case of a rule to inspect using a known server Certificate and Key, there is no option to specify Certificate Status Actions; the override setting and segment default actions have no effect.

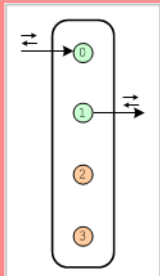
The Plaintext Marker panel controls how the generated flow with the decrypted payload is marked or if it is marked at all. The options are to have these flows be marked with:

- a VLAN tag – the VLAN ID used is configurable
- a modified source MAC address
- no marking

Because this example is a passive-tap segment, all three options are available. In the case of an active-inline segment, the “no marking option” is not available. Generated flows must be marked in order for the SSL appliance to identify them when they returned to it by the attached security appliance. In the example shown in the previous graphic, the generated flows are sent out with no marking.

The Interface column in the Segment does not shows interface numbers. These are allocated when the segment is activated. Click **Activate** in the tool block at the top right of the segment panel to select the segment you want to activate. During the activation process you can select the ports that for the segment, any copy ports, and the modes for the copy ports. Click **Apply** on the bottom left to apply the controls. The following graphic shows these interface numbers which indicate how the device is wired up to the network.

Passive-Tap Segment Activated

Segments					
Mode of Operation	Segment ID	Ruleset	Interfaces	Session Log	Comment
	A	passive-ruleset	1, 2	Enabled	

In this example:

- Port 1 connects to the network tap device that is feeding traffic to the SSL appliance
- Port 2 connects to the first passive security appliance

The red background indicates that this segment is not activated. See [Monitoring the System](#) on page 107 for details on the session log and other monitoring tools.

Example Passive-Inline Mode Inspection

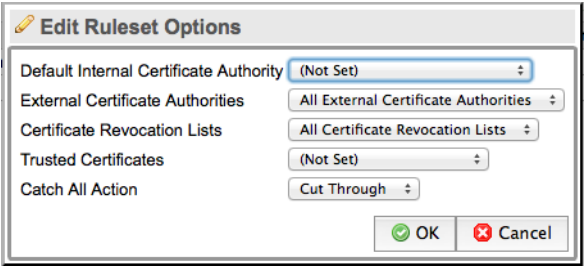
The following example shows the steps needed to configure the SSL appliance to inspect traffic that is intended for a number of SSL servers where you cannot obtain a copy of the private key and certificate. In this example, the SSL appliance is deployed in passive-inline mode as described in [Passive-Inline Mode](#) on page 32. This example illustrates the use of certificate re-sign to inspect traffic, how to use custom lists to enable a single rule to apply to traffic going to multiple destinations, and how to apply policy to SSL traffic that is not being inspected. The Internal CA used in this example is shown in previous section.

The steps involved are:

1. Create or load an Internal CA certificate and key into the SSL appliance (see [Creating a CA](#) on page 83).
2. Create a ruleset that contains rules to inspect traffic going to specific destinations.
3. Create a list of destinations for use by a single rule.
4. Create a segment for passive-inline operation.
5. Activate the segment to start inspection.

The following graphic shows the edit options screen for a ruleset called `passive-inline-example` that has already been added to the rulesets on the system. The internal CA created above is selected as the default Internal Certificate Authority.

Passive-Inline Ruleset Creation



Edit Ruleset Options	
Default Internal Certificate Authority	(Not Set)
External Certificate Authorities	All External Certificate Authorities
Certificate Revocation Lists	All Certificate Revocation Lists
Trusted Certificates	(Not Set)
Catch All Action	Cut Through
<div>OK Cancel</div>	

Before adding rules to this ruleset, we will create a list of Common Names (CN) that will allow a single rule to apply to SSL sessions to multiple destinations.

The following graphic shows the list that we are going to use in this example.

List of Distinguished Names

Distinguished Names Lists

Name
sslmg-unsupported-sites
sample-names
webmail

Distinguished Names

Item
account.google.com
mail.yahoo.com

Click on the icon in the Distinguished Names Lists area and give the new list a name (in this example, **webmail**). Select the new empty list in the Distinguished Names Lists area and click the plus icon to be added to the list. In this example, two Common Names have been added to the list. Click **Apply** at the bottom of the screen to complete the process and save the new list to disk.

You can now go to the ruleset and add a rule to use this list. The following graphic shows the list creation box with the relevant parameters configured.

Rule to Inspect Using Certificate Re-Sign a DN List

Insert Rule

Action: Decrypt (Resign Certificate)

Comment:

Internal CA: Sourcefire, IT, demonstration

Cipher Suite List: (Not Set)

☐ Trusted Certificate
☒ Trusted Certificates (Not Set)
☐ Subject DN
☒ Subject DN List: webmail
☒ Issuer DN
☐ Issuer DN List: (Not Set)
☒ Source IP
☐ Source IP List: (Not Set)
☒ Destination IP
☐ Destination IP List: (Not Set)

Destination Port: 443

Certificate Status:

- revoked
- self-signed
- valid
- invalid-signature
- expired
- invalid-issuer
- not-valid-yet

OK Cancel

Note that the radio button beside **Subject DN List** is checked and **webmail** has been selected from the drop-down menu. In this example, we have also configured the **Destination Port** to be 443. This rule will inspect any traffic going to a server that has a CN that is in the **webmail** list where the destination port number is 443. If there was any traffic to one of the servers on the list that had a destination port number other than port 443, this rule would not be triggered.

IMPORTANT! In this example, the entries added to the list were all Common Names and entered into list input box. You can include other elements of the x509 certificate in a list by specifying what the item is when it is added. The default element type is a Common Name. More details on how to include other elements of the x509 certificate in a list are given later in this document.

The default action for this ruleset is *cut through*. Any SSL traffic which does not match the rule will be cut through and is not inspected. If we wanted to prevent traffic to a specific SSL site, you can add another rule could be ruleset that matches on the specific Common Name for that site and drops the traffic.

The following graphic shows how the ruleset appears after a second rule has been added to decrypt traffic going to 10. 100. 100. 54.

Passive-Inline Ruleset with Two Rules Defined

Rulesets

Name	Rule Count
active-inline-ruleset	1
passive-inline-ruleset	2
passive-ruleset	1

Ruleset Options

Default Internal Certificate Authority: (Not Set)
External Certificate Authorities: All External Certificate Authorities
Certificate Revocation Lists: All Certificate Revocation Lists
Trusted Certificates: (Not Set)
Catch All Action: Cut Through

Rules

Match Fields	Action	Comment
known-certificates-with-keys[all-known-certificates-with-keys]	Decrypt (Certificate and Key known)	
known-keys[all-known-keys],dst-ip[10.100.100.54]	Decrypt (Key known)	

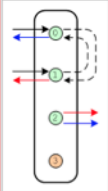
Click **Apply** at the bottom of the screen and you will be able to see that the rules are now part of the ruleset.

The final part of the process is to create a segment, configure it to use the ruleset just created, and activate it.

To create a Segment, select the **Policies / Segments** menu option. Click the plus button in the Segments table and follow the same process as in the earlier example, and choose a passive-inline segment type. In the Policy Changes

notification, click **Apply** to save the CA to disk. The following graphic shows the segment after it has been completed, saved and activated.

Passive-Inline Segment Configuration

Mode of Operation	Segment ID	Ruleset	Interfaces	Session Log	Comment
	A	passive-inline-example	1, 2, 3	Enabled	

Undecryptable Actions	
Compression:	Cut Through
SSL2:	Cut Through
TLS 1.2:	Cut Through
Diffie-Hellman in Passive-Tap mode:	Cut Through
Client Certificate:	Reject
Cipher Suite (including Export):	Cut Through
Uncached:	Cut Through

Certificate Status Actions	
Invalid Issuer:	(Not Set)
Invalid Signature:	(Not Set)
Expired:	(Not Set)
Not Valid Yet:	(Not Set)
Self Signed:	(Not Set)
Revoked:	(Not Set)
Status Override Order:	Rule over Segment

Plaintext Marker	
Type:	(Not Set)

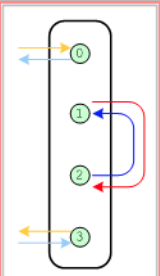
Failure Mode Options	
Software Failure Action:	Fail-to-wire (Auto Recovery)
High Availability:	Disabled

Notice that:

- The ruleset created above is configured as the ruleset to be used for this segment.
- The session log has been turned on for this segment.
- Interfaces 1, 2, 3 and 4 are allocated to this segment; interface 4 is not used.

The following graphic shows the active segment status and the interface numbers which indicate how the device is wired to the network.

Passive-Inline Segment Active

Segments					
Mode of Operation	Segment ID	Ruleset	Interfaces	Session Log	Comment
	A	active-inline-ruleset	1, 2, 3, 4	Enabled	

In this example:

- Interfaces 1 and 2 connect to the network, making the SSL appliance a bump-in-the-wire
- Interface 3 connects to the attached passive security appliance

A green background indicates that the segment is active. The SSL session log and SSL statistics screens display any SSL traffic. See [Monitoring the System](#) on page 107 for details on the session log and other monitoring tools.

Example Active-Inline Mode Inspection

The following example shows the steps needed to configure the SSL appliance to inspect traffic and to pass the inspected traffic through an active-inline security appliance. In this example, the SSL appliance is deployed in active-inline mode as described in [Active-Inline Mode](#) on page 33. This example illustrates the use of both certificate re-sign and known server key mechanisms to inspect traffic. It also illustrates the use of custom lists and how to apply policy to SSL traffic that is not being inspected.

The steps involved are:


1. Create or load an Internal CA certificate and key into the SSL appliance.
2. Load one or more server certificates and keys into the SSL appliance.
3. Create a ruleset that contains rules to inspect traffic going to specific destinations.
4. Create a list of destinations for use by a single rule.
5. Create a list of local servers and which keys/certificates are available.
6. Create a segment for active-inline operation.
7. Activate the segment to start inspection.

The only steps in this process that have not already been covered in earlier examples are:

- creation of a list of known server keys and certificates
- creation of a ruleset that includes both known server key inspection and certificate re-sign inspection
- creation of an inline-active segment

These steps are shown below.

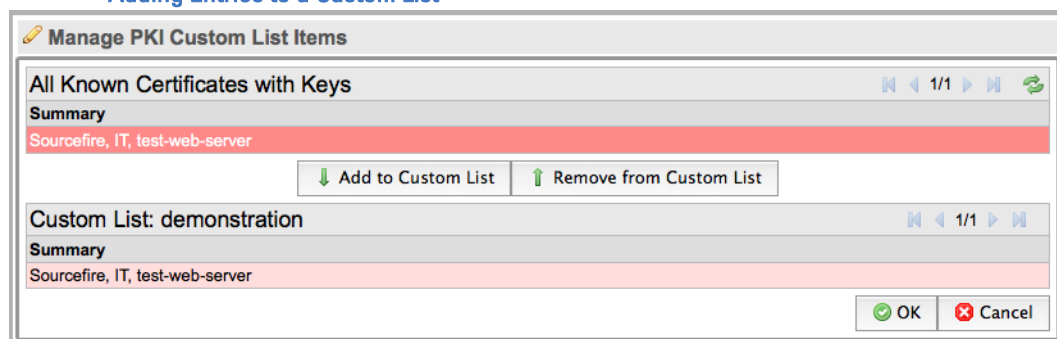
Creation of a Custom List of Known Server Keys/Certificates



This graphic shows the Known Certificates with Keys List box after a list called **demonstration** has been added and saved. To add entries to the list, highlight the **demonstration** list and click on the icon in the Known Certificate with Keys section.

To add keys and certificates to the custom list, copy the keys and certificates from the **all-known-certificates-with-keys** list. The following graphic shows the mechanism used to copy the desired keys/certificates to the custom list.

Adding Entries to a Custom List



The top section of the box lists all the keys/certificates that are present in the **all-known-certificates-with-keys** list. Highlight and click the **Add to Custom List** button to copy the item into the customer list. In the above graphic, the key/certificate for **test-web-server** has been copied. After the keys/certificates are copied, click **OK**. Click **Apply** in the Policy Changes notification block to complete the process and to save the CA to disk.

The ruleset for this example is shown in the following graphic and includes three rules.

Active-Inline Ruleset

Rulesets

Name	Rule Count
active-inline-ruleset	1
passive-inline-ruleset	2
passive-ruleset	1

Ruleset Options

Default Internal Certificate Authority:	(Not Set)
External Certificate Authorities:	All External Certificate Authorities
Certificate Revocation Lists:	All Certificate Revocation Lists
Trusted Certificates:	(Not Set)
Catch All Action:	Cut Through

Rules

Match Fields	Action	Comment
known-certificates-with-keys[all-known-certificates-with-keys]	Decrypt (Certificate and Key known)	
known-keys[all-known-keys],dst-ip[10.100.100.54]	Decrypt (Key known)	

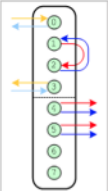
The first rule uses the [local -servers](#) list to inspect traffic using known server key/certificate mechanisms. The second rule rejects any SSL sessions that have an expired server certificate. The third rule uses the webmail destinations list to inspect traffic using certificate re-sign.

WARNING! Rules are processed sequentially from top to bottom. The up and down arrow buttons can be used to alter the position of a rule in the Rules block.

The final part of the process is create a segment, configure it to use the ruleset above and then to activate it. To create a Segment, go to the [Policies / Segments](#) menu option, and click [Add](#) in the Segments table.

The following graphic shows the segment configuration after it has been saved and activated.

Active-Inline Segment Configuration

Mode of Operation	Segment ID	Ruleset	Interfaces	Session Log	Comment
	A	active-inline-example	1, 2, 3, 4, 5, 6	Enabled	

Undecryptable Actions	
Compression:	Cut Through
SSL2:	Cut Through
TLS 1.2:	Cut Through
Diffie-Hellman in Passive-Tap mode:	Cut Through
Client Certificate:	Reject
Cipher Suite (including Export):	Cut Through
Uncached:	Cut Through

Certificate Status Actions	
Invalid Issuer:	(Not Set)
Invalid Signature:	(Not Set)
Expired:	(Not Set)
Not Valid Yet:	(Not Set)
Self Signed:	(Not Set)
Revoked:	(Not Set)
Status Override Order:	Rule over Segment

Plaintext Marker	
Type:	Source MAC
MAC Address:	00:15:4D:00:00:D5

Failure Mode Options	
Software Failure Action:	Fail-to-wire (Auto Recovery)
High Availability:	Disabled

In this example you can see:

- The configuration allows the connection of an active security appliance, such as an IPS.
- The configuration allows the connection of two passive security appliances, each of which receive a copy of the traffic being sent to the active appliance.
- The session log is enabled for this segment.
- Generated flows containing decrypted traffic are marked by changing the source MAC address to the value indicated.

Chapter 5

Web-Based Management Interface (WebUI)

This chapter provides details of all the facilities provided by the WebUI on the SSL appliance. Each top level menu option is covered by a specific section that details all the features available and how they are used.

To connect to the web interface on the SSL appliance, start a web browser and enter the host name or IP address of the appliance in the address bar. To view the current IP address and host name of the appliance on the front panel LCD screen, press the bottom right button on the keypad. If the host name has not been set yet, or if the host name does not map to the IP address, the IP address must be used.

Browser Configuration

If you attempt to access the web interface without the correct certificate installed in the web browser, the browser will display a warning dialog box or message. This is the normal and correct behavior for the web browser. To prevent the warning message being displayed, the browser must be configured to trust the certificate being used by the webserver in the SSL1500, SSL2000, or SSL8200.

There are two ways that the browser can be made to trust the SSL appliance certificate. An SSL server certificate issued by a trusted CA can be loaded into the SSL appliance, this will be used by the internal web server; or you can configure the browser to trust the *self-signed* server certificate that the SSL appliance uses by default.

Details on how to import an SSL server certificate to the SSL appliance are given in [Import UI Certificate/Key](#) on page 137.

If the browser generates warnings, consult your browser documentation for instructions on how to add the SSL appliance certificate to the set of trusted certificates stored in the browser.

For Chrome, clicking **Proceed anyway** allows the browser to connect to the SSL appliance.


For Firefox, clicking on the **I understand the risks** button allows access to screens that allow the certificate from the SSL appliance to be added to the set of trusted certificates within Firefox.

Login Process

The SSL appliance does not have a default username and password when it is shipped from the factory. During the initial bootstrap configuration, you can create a user name and password and log on to the system after the bootstrap phase is complete. See [Installation Process](#) on page 73 for details of the bootstrap process. Additional user names and passwords can be created on the system using the WebUI. Multiple users can be logged on to the system at the same time. The system will rate-limit login attempts to prevent attacks. The system will also timeout a session and then prompt the user for their password before allowing access again.

The following graphic shows the standard login box presented by the WebUI.

Login Box



The screenshot shows a login box with the Sourcefire logo at the top. Below the logo are two input fields: "User ID" and "Password". At the bottom of the box is a "Login" button with a user icon.

Screen Layout Explained

The management interface screens display different types of information in specific areas on the screen. The basic organization of the management screens is described below.

The following graphic shows information that is present at the top and bottom of every screen.

Management Screen Basic Layout




The top of the screen contains five menus, a refresh button and, when a refresh is occurring, a spinner to indicate this fact. The five menu items are explained in detail in later sections. The bottom of the screen shows a status bar that is always present and which shows the following information:

- Current date in the format YYYY-MM-DD
- Copyright notice
- SSL appliance Model Number (SSL Appliance 2000 or SSL Appliance 8200)
- Software version currently running on the system
- Icons showing current status for: System, Load and Network
 - indicates an error
 - indicates a warning
 - indicates everything is fine

The screen displays additional information between the top and bottom bars, organized into panels. Each panel has a title bar at the top and a set of tool icons at the right hand side. The rest of the panel displays information. The set of tools available varies by panel and some of the tools may be unavailable and grayed out depending on how the panel is being used. Panels may also be empty, in which case only the title bar will be visible.



A panel that only displays information has the refresh tool icon at the right side of the title bar. Click the refresh tool to refresh the data in the panel. The following graphic shows an example of a display only panel.

Example Information Display Panel

Chassis FRU Info 	
Chassis Part Number	Chassis Serial Number
NFPP-1U-AC	515-11012200400025




Panels that contain editable data have an edit tool icon in addition to the refresh tool icon. The following graphic is an example of a panel that displays configuration data and allows it to be edited.




Example Configuration Edit Panel

Management Network  	
MAC Address:	00:15:4d:00:85:c5
MTU:	1500
DHCP:	True
IP Address:	192.168.3.214
Netmask:	255.255.255.0
Default Gateway:	192.168.3.1
Hostname:	localhost.localdomain
Primary Nameserver:	
Secondary Nameserver:	




Panels can also be linked to other panels so that an action taken in one panel affects the related panel. The following graphic shows an example of two linked panels.

Example of Linked Panels




Distinguished Names Lists   	
Name	
sslmg-unsupported-sites	
sample-names	
webmail	

Distinguished Names   	
Item	
account.google.com	
mail.yahoo.com	





The top Distinguished Names Lists panel contains details of lists that are stored in the system and has tool icons allowing the following actions (in addition to the refresh action and multi-page tools):

- Addition of a new list 
- Deletion of an existing list 
- Cloning of an existing list 

When a row in the top Distinguished Names Lists panel is selected the lower Distinguished Names panel will show the names contained in the list that has been selected and provides tool icons that allow:

- Addition of a name 
- Editing of a name  (grayed out if a name has not been selected)
- Deletion of a name  (grayed out unless a name has not been selected)

Some panels indicate which page the panel is currently displaying, and tool icons that allow movement between pages within the panel. The tool icons are explained below.

- Jump to first page 
- Jump to last page 
- Move forward one page 
- Move backward one page 

You can move directly to a particular page by clicking on the numbers between the tool icons and typing in the number of the required page.

IMPORTANT! Multi-page panels are configured to display a maximum number of rows so the maximum number of pages that the panel supports is determined by the page size that is configured. They also have a built in multiplier that is used in conjunction with the number of rows value. See [Preferences](#) on page 139 for more details. As an example, the SSL Session log holds 1024 entries which with the default row setting of 10 will mean there is a maximum of 64 pages. The SSL Statistics panel has a multiplier of 1.6 so with the default row setting of 10 this will mean there are 16 rows displayed in the SSL statistics panel. If the default row count was set to 20 then the SSL Statistics panel would have 32 rows.

Monitoring the System

The Monitor menu contains eight options that provide details on the operation of the system and allow the collection of diagnostic and debug information.

Hover over the Monitor Menu title to display the menu options in the following graphic. These options are described in detail below in the order in which they appear on the menu.

Monitor Menu Options

Monitor	Policy
Dashboard	
System Log	
SSL Session Log	
SSL Statistics	
Certificates	
Errors	
Diagnostics	
Debug	

Dashboard

The dashboard display contains seven panels, described below.

The following graphic shows the segment status panel which displays the status of currently active segments.

Dashboard Segment Status Panel

Segments Status				
Segment ID	Interfaces Used	Interfaces Down	Main Mode	Failures
A	1, 2		Passive Tap	

The Segment ID is a unique identifier that distinguishes this segment from other segments in the system. The Interface numbers identify the physical ports used by this segment. If any of the interfaces used by the segment are down then the interface numbers will show in the Interfaces Down column. Main Mode indicates the operating mode of the segment and the Failures column records failure details. The Manually Unfail icon is normally available if the segment is in a failure mode that requires manual intervention to clear the failure.

The following graphic shows the Network Interfaces panel which contains a row for every interface that is installed in the system. The maximum number of rows in an SSL Appliance 2000 is 12, and 16 for an SSL Appliance 8200. The link state column will show the link speed for 1G NetMods.


Dashboard Network Interfaces

Network Interfaces					
Port	Type	Link State	RX Packets/Bytes	TX Packets/Bytes	RX Drops
1	1G	1G	776/116320	0/0	0
2	1G	1G	0/0	776/116320	0
3	1G	Down	0/0	0/0	0
4	1G	Down	0/0	0/0	0

Each row shows the interface type, speed (10Mbps, 100Mbps, or GigE), transmit, and receive statistics. The only tool provided for this panel is the refresh button.


The following graphic shows the current CPU utilization as a percentage of the total capacity of the CPU. The only tool provided for this panel is the refresh button.

Dashboard CPU Load %

CPU Load %																	
cpu	cpu0	cpu1	cpu2	cpu3	cpu4	cpu5	cpu6	cpu7	cpu8	cpu9	cpu10	cpu11	cpu12	cpu13	cpu14	cpu15	
3.1	40.7	0	0	1	0	0	0	0	1	0	0	0	0	0	0	1	


The following graphic shows the Fan Speed panel which has the current speed values for the various fans in the system. The only tool provided for this panel is the refresh button.

Dashboard Fan Speed (RPM)

Fan Speed (RPM)							
Fan Mod 1 Inlet	Fan Mod 1 Outlet	Fan Mod 2 Inlet	Fan Mod 2 Outlet	Fan Mods 3 to 5	Left Power Supply Fan	Right Power Supply Fan	
18226	16148	18226	16550	7304	5882	16129	


The following graphic shows the Temperatures panel which includes details of temperatures and thermal margins for components within the system. The only tool provided for this panel is the refresh button.

Dashboard Temperatures (Degrees °C)

Temperatures (Degrees °C)											
Baseboard Temp	Front Panel Temp	IOH Therm Margin	Mem P1 Thrm Mrgn	Mem P2 Thrm Mrgn	P1 Therm Margin	P2 Therm Margin	NFP0 Temp	Left Power Supply Temp	Right Power Supply Temp		
35	29	-37	-28	-44	-54	-45	57	39	44		


The following graphic shows the Utilization panel which shows the percentage utilization of system memory and disk space. The only tool provided for this panel is the refresh button.

Dashboard Utilization %

Utilization %		
Disk	Memory	
2.0	24.3	

The following graphic shows the System Log panel that contains the most recently generated system log entries. This panel automatically refreshes.

Dashboard System Log

System Log			
Time	Process	Log	
Jul 24 15:17:46	sslmange[4831]	Store update detected: Policy	
Jul 24 15:17:46	ssldata[4834]	Activation request received. Activation pending	
Jul 24 15:17:46	sslmange[4831]	Activation request sent to data-plane	

System Log

The System Log screen contains a single multi-page panel enabling all entries in the system log to be viewed. The panel has multi-page navigation buttons, the refresh button, and a search button. Click on the **Search** button to bring up a dialog box to filter log entries.

System Log Panel

System Log				
Time	Process	Log		
Dec 2 12:16:10	kernel	imklog 4.6.4, log source = /proc/kmsg started.		
Dec 2 12:16:10	rsyslogd	[origin software="rsyslogd" swVersion="4.6.4" x-pid="604" x-info="http://www.rsyslog.com"] (re)start		
Dec 2 12:16:10	rsyslogd-2039	Could no open output pipe '/dev/xconsole' [try http://www.rsyslog.com/e/2039]		
Dec 2 12:16:10	kernel	[0.000000] Initializing cgroup subsys cpuset		



To cancel a filter, open the Filter on Process input box, delete the text in the input field, and click **OK**.



SSL Session Log

The SSL Session Log screen contains a single multi-page panel enabling all entries in the last 64 pages of the SSL Session log to be viewed.

Session Log Panel

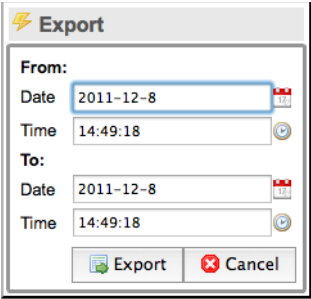
SSL Session Log										
Start Time	Segment ID	SrcIP:Port	DstIP:Port	Subject	Certificate Status	Cipher Suite	Action	Status		
Dec 08 14:49:17.406 *	A	192.168.0.121:49217	192.168.0.80:443	--unknown--		TLS_RSA_WITH_AES_128_CBC_SHA	Cut Through	Master key invalid		
Dec 08 14:49:17.406	A	192.168.0.121:49217	192.168.0.80:443	--unknown--		TLS_RSA_WITH_AES_128_CBC_SHA	Invalid	Master key invalid		
Dec 08 14:49:17.406 *	A	192.168.0.121:49215	192.168.0.80:443	--unknown--		TLS_RSA_WITH_AES_128_CBC_SHA	Cut Through	Master key invalid		

The panel has multi-page navigation buttons, a refresh button, a **View Details** button, and an **Export** button and two filter buttons: filter on error: , and no filter . Click the **Export** button to view the range of SSL session log entries that can be exported.

The filter on error button  causes the session log to only display entries for flows with an inspection error. The no filter button  causes the session log to display all entries.

The following graphic shows the export dialog box that allows the start and end date and time that the exported session logs should cover.

Session Log Export Box


A screenshot of the 'Export' dialog box. The dialog has a title bar with a lightning bolt icon and the word 'Export'. It contains two sections: 'From:' and 'To:'. Each section has a 'Date' field with a calendar icon and a 'Time' field with a clock icon. Both 'From' and 'To' fields are currently set to '2011-12-8' and '14:49:18' respectively. At the bottom, there are two buttons: 'Export' with a green document icon and 'Cancel' with a red X icon.

Click [Export](#) button to save the export file to the default location for your browser or a user-specified location. The saved file contains a set of `.bin` files and a file that contains the public certificates used in the SSL sessions captured in the session log. To view the session log data, the `.bin` files must be processed with a tool to extract the data in a user readable form. The tool and documentation for the tool are provided separately.

The Session Log includes the following details for each SSL session recorded in the log:

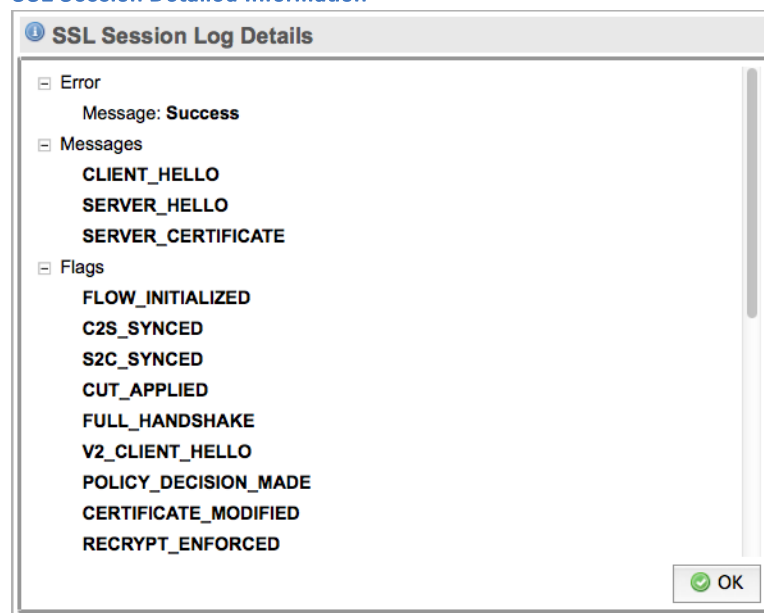
- Start date and time
- Segment ID for the segment the SSL session occurred on
- IP source and destination address and port number
- Subject details from the server certificate used during the session
- Status of the server certificate
- Cipher Suite that was used for the session
- Action taken by the SSL appliance for this session
- Status for the session

Entries in the session log are ordered from most recent to oldest. The first row on page 1 is the most recent entry and the last row on page 64 is the oldest entry.

The [View Details](#) button  is only active when a row in the SSL Session Log panel has been selected. Clicking on the [View Details](#) button will open up a dialog box

showing more details about the selected session. The following graphic shows an example of the detail available for a successful session.

SSL Session Detailed Information



SSL Statistics

The SSL Session Log screen contains a single multi-page panel enabling all entries in the last 64 pages of the SSL Statistics log to be viewed. The panel has multi-page navigation buttons, and a refresh button.

The following graphic displays an example of available statistics information.

SSL Statistics

SSL Statistics								
Timestamp	#Detected	#Done	#Ignored	#Decrypt	#Decrypt Done	#Error	Detected	Decrypt
Dec 08 16:13:19	16	16	12	0	0	16	0	0
Dec 08 16:13:18	16	16	12	0	0	16	0	0
Dec 08 16:13:17	16	16	12	0	0	16	0	0
Dec 08 16:13:16	16	16	12	0	0	16	0	0
Dec 08 16:13:15	16	16	12	0	0	16	0	0
Dec 08 16:13:14	16	16	12	0	0	16	0	0
Dec 08 16:13:13	16	16	12	0	0	16	0	0
Dec 08 16:13:12	16	16	12	0	0	16	0	0
Dec 08 16:13:11	16	16	12	0	0	16	0	0
Dec 08 16:13:10	16	16	12	0	0	16	0	0
Dec 08 16:13:09	16	16	12	0	0	16	0	0
Dec 08 16:13:08	16	16	12	0	0	16	0	0
Dec 08 16:13:07	16	16	12	0	0	16	0	0
Dec 08 16:13:06	16	16	12	0	0	16	0	0
Dec 08 16:13:05	16	16	12	0	0	16	0	0
Dec 08 16:13:04	16	16	12	0	0	16	0	0

Statistics are collected every second and each row in the table holds the data for a collection interval. All the counts are cumulative except for the Detected and Decrypted columns. The Detected and Decrypted columns show the instantaneous number of sessions in each category at the point the data was collected, this is not the total number of sessions that may have been in that category over the one second period.

Entries in the Statistics panel are ordered from most recent to oldest. The first row on page 1 is the most recent entry and the last row on page 64 is the oldest entry.

Invalid Certificates

The Invalid Certificates screen contains a single tabbed panel that shows details of invalid certificates received by the SSL appliance. The panel has an acknowledge tool, a refresh button, and an export button. The export button allows details of all invalid certificates to be exported to a .csv file. The tabs show details for different types of invalid certificate state.

The following graphic displays details of certificates which were issued by invalid (untrusted) Certificate Authorities.

Invalid Certificates Panel

Invalid Issuer	Invalid Signature	Revoked	Self-Signed	Acknowledged
Certificates with Invalid Issuer				
First Seen on	Segment ID	Certificate	Details	

By clicking the relevant tab you can view details for other types of invalid certificates. The following graphic shows details of self-signed certificates.

Invalid Certificates Panel with Self-Signed Certificates Details

Invalid Issuer	Invalid Signature	Revoked	Self-Signed	Acknowledged
Self-Signed Certificates				
First Seen on	Segment ID	Certificate		

If a certificate is invalid for more than one reason, it will appear on more than one tab. The acknowledge tool can be used to notify the system that the certificate status has been acknowledged. After a certificate has been acknowledged, it appears on the acknowledged tab only. To acknowledge a certificate, select the certificate and click on the check mark icon. Acknowledged certificates are not included in details on invalid certificates that are collected in the system log files.

IMPORTANT! Invalid certificate details are automatically cleared from any tab when the segment that they occurred on is deactivated.

Errors

The Errors screen contains a single panel that shows SSL Error counts for each active segment. The panel has multi-page controls, a refresh button, and an **Export** button. The **Export** button allows details of all errors to be exported to a .csv file.

The following graphic displays one entry indicating the number of undecryptable sessions that have occurred on segment A.

SSL Error Counts Panel

SSL Error Counts			
Segment ID	Code	Message	Count
A	0x32001442	Undecryptable	8

There may be multiple rows for a single segment if there have been more than one type of error seen on that segment. Whenever a segment is activated or deactivated the error counts associated with that segments are reset to zero.

Diagnostics

The Diagnostics screen contains a dialog box that allows the user to specify what types of information are included in the diagnostic file and cause the file to be generated.

The following graphic shows the dialog box with data for Policy, Platform, and SSL statistics selected for the diagnostic file.

Diagnostics Box

The image shows a 'Diagnostics' dialog box with a title bar and a pencil icon. Inside, there is a section 'Generate diagnostics for:' with several checkboxes. The checkboxes for 'Policy', 'Platform', and 'SSL Statistics' are checked, while 'PKI', 'Host Statistics', and 'NFP Statistics' are unchecked. Below the checkboxes are two date pickers: 'Start' with the date '2011-12-1' and 'End' with the date '2011-12-8'. At the bottom are 'OK' and 'Cancel' buttons.

Click **OK** to create the file.

WARNING! Including the SSL Statistics, the Host Statistics, or the NFP statistics may result in a large diagnostic file. These should only be included if required.

Debug

The Debug display contains a multi-page network statistics panel intended to assist with debugging issues with the SSL appliance. Support personnel may ask for information from the debug screens. The NFE Network Statistics panel contains information that may be useful to a user in diagnosing configuration issues.

The panel has multi-page navigation buttons and a [Refresh](#) button.

The following graphics show the SSL debug statistics panel.

Debug SSL Statistics

L_ssl_decrypt_busy	0
L_ssl_detected	0
L_ssl_intercepted	0
M_handshake	0
M_record_tmp	0
M_session_cache	0
T_modmath_dh_replies	0
T_modmath_dh_requests	0
T_modmath_epms_replies	0
T_modmath_epms_requests	0
T_modmath_errors	0
T_modmath_x509_replies	0
T_modmath_x509_requests	0
T_ssl_alert_c_bad_record_mac	0
T_ssl_alert_c_close_notify	0
T_ssl_alert_c_decrypt_error	0
T_ssl_alert_c_fatal	0

T_ssl_full_handshake_intercepted	0
T_ssl_handshake_error	16
T_ssl_ignored	12
T_ssl_intercept_done	0
T_ssl_intercepted	0
T_ssl_reused_session_ignored	12
T_ssl_reused_session_intercepted	0

Counts with names that begin with **L_** are instantaneous values. Counts that begin with **T_** are cumulative totals. If SSL traffic is passing through the SSL appliance, some counts should increment. Counts with **modmath** in the name relate to PKI activity during the SSL handshake phase. Use the set of counters on page 2 to determine if SSL traffic is present and how it is being dealt with. The reused session counts help explain why some SSL flows may not be being inspected.

The NFE Network Statistics panel shows details of traffic to and from the Network Flow Engine (NFE) acceleration card used in the SSL appliance. The NFE card has two 10Gbps links that connect to an ethernet switch which in turn connects to the set of NetMods that provide the external interfaces on the SSL appliance.

The following graphics show details for four NFE links. Because the SSL Appliance 2000 has two NFE links, two of the columns will always show zero

counts. The SSL Appliance 8200 has four NFE links and will display the two remaining columns.

Debug NFE Network Statistics

NFE Network Statistics		
	nfe 0	nfe 1
BadCRC	0	0
BadOctetsReceived	0	0
BroadcastFramesReceived	160	0
BroadcastFramesSent	160	0
CRCErrorsSent	0	0
Collisions	0	0
ControlFrameReceived	0	0
ExcessiveCollisions	0	0
FCReceived	0	0
FCSent	157810061	157810060
Fragments	0	0
Frames1024toMaxOctets	0	0
Frames128to255Octets	64	0
Frames256to511Octets	176	0
Frames512to1023Octets	0	0
Frames64Octets	0	0

NFE Network Statistics		
	nfe 0	nfe 1
Frames65to127Octets	520	16
GoodOctetsReceived	60688	576
GoodOctetsSent	60688	576
GoodUnicastFramesReceived	220	8
InRangeLengthErrorsReceived	0	0
Jabber	0	0
LateCollisions	0	0
MulticastFramesReceived	0	0
MulticastFramesSent	0	0
OutOfRangeLengthErrorsReceived	0	0
Oversize	0	0
ReceiveFIFOOverrun	0	0
RxErrorFrameReceived	0	0
SentDeferred	0	0
SentMultiple	0	0
SymbolErrorReceived	0	0

NFE Network Statistics		
	nfe 0	nfe 1
TxBackoff	0	0
TxCARRIERSENSEERRORS	0	0
TxExcessiveDefer	0	0
Undersize	0	0
UnicastFramesSent	220	8

Configuring Segments and Policies

The **Policies** menu contains five options that allow the configuration of segments and the definition of policies and rules that determine how SSL traffic is handled and which SSL traffic is inspected.

The following graphic shows the options available on the **Policies** menu.

Policies menu Options

Policies	PKI
Rulesets	
Segments	
Distinguished Names List	
IP Address List	
Cipher Suites List	

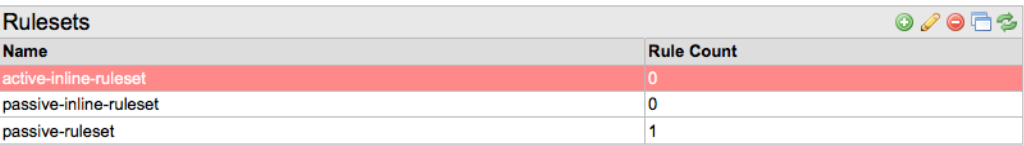
The first two options allow configuration of Rulesets and Segments; the remaining options allow configuration of lists that can be used within Rulesets. These options are described in detail below in the order in which they appear on the menu.

Rulesets

The Rulesets display contains three panels with the lower two panels displaying information that varies depending on the row selected in the first panel. These panels are described below. Rulesets contain the rules and policies that control how SSL traffic is handled and are associated with one or more segments. Rulesets can also exist that are currently not associated with any segment.

The following graphic shows the Rulesets panel with three existing rulesets being shown.

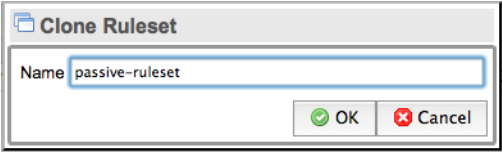
Rulesets Box

A screenshot of the 'Rulesets' panel in the WebUI. It features a table with two columns: 'Name' and 'Rule Count'. The table lists three rulesets: 'active-inline-ruleset' with 0 rules, 'passive-inline-ruleset' with 0 rules, and 'passive-ruleset' with 1 rule. The 'active-inline-ruleset' row is highlighted in red. Above the table, there are icons for adding, editing, deleting, and cloning rulesets.

Name	Rule Count
active-inline-ruleset	0
passive-inline-ruleset	0
passive-ruleset	1

Each existing ruleset occupies one row in the table and the right hand column shows the number of rules that are currently within that ruleset. Tools icons on this panel allow addition or removal of a ruleset or cloning of a ruleset. The remove and cloning icons will be grayed out unless an entry in the table is selected. If the clone tool is used a dialog box appears to allow the rulesets name for the clone to be input. The following graphic shows the dialog box. A similar dialog box will appear if the add ruleset option is selected.

Rulesets Clone Box


A screenshot of the 'Clone Ruleset' dialog box. It has a title bar with a folder icon and the text 'Clone Ruleset'. Inside, there is a text field labeled 'Name' with the value 'passive-ruleset'. At the bottom right, there are two buttons: 'OK' with a green checkmark icon and 'Cancel' with a red X icon.

Clone Ruleset
Name: passive-ruleset
OK Cancel

To cause the second and third panels to display information a ruleset entry in the Rulesets panel must be selected. This is done by clicking on an entry, which will highlight the entry in the Rulesets panel and cause the Rulesets Options panel to expand. The panel becomes active, and the Rules panel displays the rules that exist within the selected ruleset.

The following graphic shows the Ruleset Option panel that allows configuration of settings that apply to the ruleset. The panel provides an edit tool in addition to the refresh button.

Ruleset Option Panel

A screenshot of the 'Ruleset Options' panel. It has a title bar with a pencil icon and a refresh icon. The panel lists several configuration options on the left and their current values on the right. The options are: 'Default Internal Certificate Authority' (Not Set), 'External Certificate Authorities' (All External Certificate Authorities), 'Certificate Revocation Lists' (All Certificate Revocation Lists), 'Trusted Certificates' (Not Set), and 'Catch All Action' (Cut Through).

Ruleset Options	
Default Internal Certificate Authority:	(Not Set)
External Certificate Authorities:	All External Certificate Authorities
Certificate Revocation Lists:	All Certificate Revocation Lists
Trusted Certificates:	(Not Set)
Catch All Action:	Cut Through

The following graphic shows the edit box with drop-down menus to allow selection of the desired settings for this ruleset.

Ruleset Options Edit Box

Edit Ruleset Options

Default Internal Certificate Authority: (Not Set)

External Certificate Authorities: All External Certificate Authorities

Certificate Revocation Lists: All Certificate Revocation Lists

Trusted Certificates: (Not Set)

Catch All Action: Cut Through

OK Cancel

The options that can be configured are:

- Default Internal Certificate Authority: used if no other internal CA is specified in a rule entry
- External Certificate Authorities: selects the list of trusted external CAs that will be checked against when SSL sessions are processed by rules within this ruleset
- Certificate Revocation Lists: selects the set of CRLs that will be checked against when SSL sessions are processed by rules within this ruleset
- Trusted Certificates: selects the set of trusted certificates that will be checked against when SSL sessions are processed by rules within this ruleset
- Catch All Action: defines what happens to an SSL session that does not trigger any rules within this ruleset

The following graphic shows the Rules panel, which displays the rules currently defined in the ruleset being edited and includes a multi-page selection tool, add/edit/delete tools, plus move up and move down tools. Use the multi-page selection tool to move between pages of rules when there are many rules in the ruleset.

Rules Panel

Rules		
Match Fields	Action	Comment
known-certificates-with-keys[all-known-certificates-with-keys]	Decrypt (Certificate and Key known)	
known-keys[all-known-keys],dst-ip[10.100.100.54]	Decrypt (Key known)	

Click on the add button to open up the Insert Rule box, shown in the following graphic.

Insert Rule Box

Insert Rule

Action: Cut Through

Comment:

Cipher Suite List: (Not Set)

☐ Trusted Certificate

☒ Trusted Certificates (Not Set)

☒ Subject DN

☐ Subject DN List (Not Set)

☒ Issuer DN

☐ Issuer DN List (Not Set)

☒ Source IP

☐ Source IP List (Not Set)

☒ Destination IP

☐ Destination IP List (Not Set)

Destination Port:

Certificate Status: revoked
self-signed
valid
invalid-signature
expired
invalid-issuer
not-valid-yet

OK Cancel

The first option in this box is a drop-down menu allowing selection of the type of rule that is to be created. Choosing an option from the drop down causes the insert rule box to change so that it only contains the fields relevant for the type of rule that has been chosen. See [Ruleset Policies](#) on page 37 for an explanation of the different parameters that can be configured for the different types of rules.

If the drop down is used to select Decrypt (Certificate and Key known), the Insert Rule box will change:

Insert Known Certificate and Rule Key

If there is more than one rule specified in a ruleset, the position of a rule in the Rules table becomes important. Rules are processed sequentially from top to bottom. If a more generic rule occurs in front of a more specific rule then the generic rule will be processed first and will always be used. For example, the following graphic shows a table with three rules:

Rules Table Showing Why Position is Important

Rules		
Match Fields	Action	Comment
known-certificates-with-keys[private-web-servers]	Decrypt (Certificate and Key known)	Inspect traffic to private web server
known-certificates[all-trusted-certificates]	Cut Through	Don't inspect traffic
known-certificates-with-keys[public-web-servers]	Decrypt (Certificate and Key known)	Inspect traffic to public web server

The comment fields indicate that two of these rules decrypt traffic to specific servers while the remaining rule causes all traffic to be cut through. The rule that is configured to cause traffic to the public web server to be inspected is never triggered. Any SSL traffic that the system detects going to the public web server will match the second rule in the table, causing all traffic to be cut through. The bottom rule in the table is never triggered. To correct this, select the last rule and reposition it above the Cut Through rule.

WARNING! If a rule does not appear to be working, always check that it is not below a more generic rule that will apply to the traffic it is intended to match.

Segments

The Segments display contains six panels with the lower four panels displaying information that varies depending on the row selected in the second panel. These panels are described below.

Multiple segments can be defined in the system and each segment may be active or inactive. The total number of active segments in a system is controlled by how many external interfaces each segment requires and how many external interfaces are supported by the system. Attempting to activate a segment when the system has insufficient free external interfaces to meet the requirements for the segment will result in an error message and the segment will remain inactive.

Deactivating an active segment releases the external interfaces used by that segment and they become available for use by other segments. The following graphic shows the first panel on the Segments screen, which enables configuration of the default action that the system should take if it is overloaded.

Segment System Options Panel




In the example shown, the action is to cut through traffic. Other options are drop or reject. This panel only has edit and refresh tool icons.

The Segments panel (second from top) contains a row for each segment configured in the system. In addition to add, edit, delete and refresh buttons it also has activate and de-activate tool icons. See [Deployment Modes](#) on page 27 for details of the modes of operation that can be selected for a segment when it is created. [Segment Policies](#) on page 36 and [Example Passive-Tap Mode Inspection](#) on page 88, [Example Passive-Inline Mode Inspection](#) on page 94, and [Example Active-Inline Mode Inspection](#) on page 98 provide examples of how to configure segments using the Segment panel.

After a segment definition exists, click on the segment definition to display information about the segment in the lower four panels. The following graphic shows the Undecryptable Actions panel, which controls how undecryptable SSL sessions on this segment are handled.

Segment Undecryptable Actions Panel



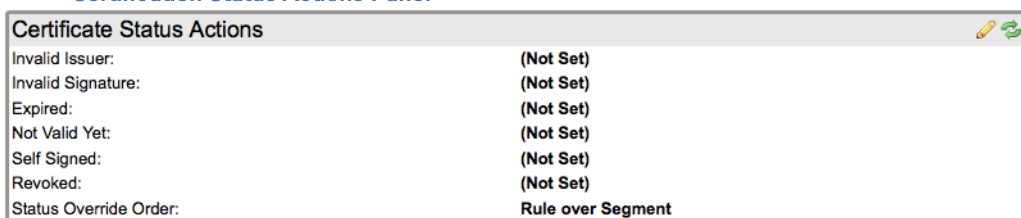
The panel has edit and refresh icons. Click the edit icon to open a dialog box with drop-down menus and select the action to be taken when a session is not decryptable for the specific reason. An SSL session cannot be decrypted for any of the following reasons:

- Compression: the system does not support inspection of SSL sessions that use compression.
- SSL2: the system only provides partial support for inspecting SSL sessions using SSLv2 (SSL v2 is an old and insecure version of SSL and its use is not recommended).

- Diffie-Hellman in passive-tap mode: when in passive-tap mode it is impossible to inspect sessions that use Diffie-Hellman (DHE) for key exchange (inspection of sessions using DHE is only possible if the inspecting device is installed inline).
- Client Certificate: the use of client certificates in some situations can prevent an SSL Session from being inspected. This action is applied when such a session is present.
- Cipher Suite: the system does not support all possible SSL cipher suites. This action is applied when an unsupported cipher suite is used by an SSL session.
- Uncached: an SSL session that is established using session re-use can be inspected only if the system has the session state for the session being re-used in its cache. This action is applied when the session state is not cached.

The following graphic shows the Certificate Status Actions panel, which controls how the system deals with SSL sessions on this segment that have particular states in the server certificate used for the session.

Certification Status Actions Panel



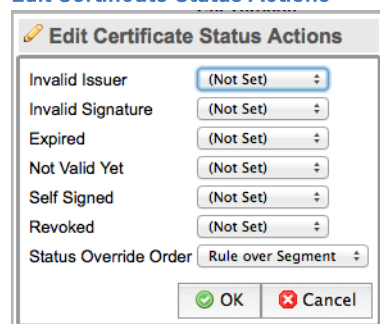
The screenshot shows a panel titled "Certificate Status Actions" with a list of certificate status categories and their corresponding actions. The categories are: Invalid Issuer, Invalid Signature, Expired, Not Valid Yet, Self Signed, Revoked, and Status Override Order. The actions for the first six categories are all set to "(Not Set)", while the Status Override Order is set to "Rule over Segment".

Category	Action
Invalid Issuer:	(Not Set)
Invalid Signature:	(Not Set)
Expired:	(Not Set)
Not Valid Yet:	(Not Set)
Self Signed:	(Not Set)
Revoked:	(Not Set)
Status Override Order:	Rule over Segment

The possible actions are, Not Set, Cut Through, Drop and Reject. Not Set means that the particular status is ignored.

The following graphic shows the Edit Certificate Status Actions and the Status Override Order.

Edit Certificate Status Actions



The screenshot shows a dialog box titled "Edit Certificate Status Actions". It contains the same list of categories as the previous panel, but each category has a dropdown menu next to it, all of which are currently set to "(Not Set)". The Status Override Order is also a dropdown menu, currently set to "Rule over Segment". At the bottom of the dialog are "OK" and "Cancel" buttons.

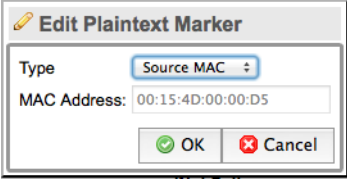
Category	Action
Invalid Issuer	(Not Set)
Invalid Signature	(Not Set)
Expired	(Not Set)
Not Valid Yet	(Not Set)
Self Signed	(Not Set)
Revoked	(Not Set)
Status Override Order	Rule over Segment

This option determines whether the segment settings in this box take precedence over any settings in rules within the ruleset used by this segment. The options are either **Rule over Segment** or **Segment over Rule**.

The Plaintext Marker panel and the Failure Mode Options panel have edit and refresh icons, and configuration of failure mode and High Availability (HA) options.

Click the edit tool for the Plaintext Market panel to produce a dialog box that controls how generated TCP flows containing inspected traffic are marked. See the following graphic.

Edit Plaintext Marker Box

A dialog box titled "Edit Plaintext Marker" with a pencil icon. It contains a "Type" dropdown menu set to "Source MAC", a "MAC Address" text field containing "00:15:4D:00:00:D5", and "OK" and "Cancel" buttons at the bottom.

There are two reasons for marking these flows:

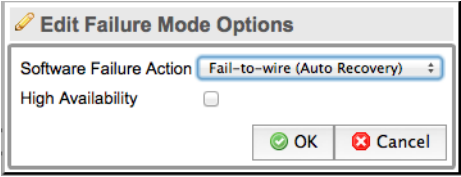
- An attached passive security appliance needs to determine which traffic it receives has been decrypted by the SSL appliance. The SSL appliance will mark all generated flows to distinguish between inspected and non-inspected traffic.
- If an SSL appliance is configured to operate in active-inline mode marking **must** be enabled to distinguish between inspected and non-inspected traffic when the traffic returns to the SSL appliance from the active security appliance.

The options available for marking generated flows are:

- Source MAC: modifies the SRC MAC address in generated flows
- VLAN: tags generated flows with a specific VLAN ID

Click on the edit icon for the Failure Mode Options panel to produce a dialog box to configure how the system deals with software failures. See the following graphic.

Segment Failure Mode Options

A dialog box titled "Edit Failure Mode Options" with a pencil icon. It contains a "Software Failure Action" dropdown menu set to "Fail-to-wire (Auto Recovery)", a "High Availability" checkbox which is unchecked, and "OK" and "Cancel" buttons at the bottom.

The following options determine how this segment will behave in the event of software failure:

- Disable Interfaces
- Drop Packets (Auto Recovery)
- Fail-to-wire (Auto Recovery)
- Fail-to-wire (Manual Reset)
- Ignore Failure

The options for High Availability Mode are:

- Disabled - HA Mode not active
- Auto Recovery - automatic recovery from failure mode when cause of failure is removed
- Manual Reset - manual action via WebUI needed to exit failure mode

Distinguished Names List

The Distinguished Names List display contains two panels; the lower panel displays information that varies depending on the row selected in the upper panel. Each DN list occupies one row in the Distinguished Names Lists panel. Tools icons on this panel allow addition or removal of a list or cloning of a list. The remove and cloning icons are grayed out unless an entry in the table is selected. By default there is a DN list called `ssl ng-unsupported-sites`, this contains the common names for sites that are known not to work if traffic to them is inspected. Selecting the list in the upper panel causes the set of names in the list to be displayed in the lower panel.

WARNING! Use a cut-through rule using the `ssl ng-unsupported-sites` list on any inline segment to enable applications using these sites to function normally.


The following graphic shows the set of names in the default `ssl ng-unsupported-sites` list.

Distinguished Names List for Unsupported Sites

Distinguished Names	
Item	
cn=account.live.com	
cn=*.update.microsoft.com	
cn=*.itunes.apple.com	
cn=*.citrixonline.com	
cn=*.logmein.com	
cn=*.mozilla.org	

All the entries contain a Common Name value, which is the CN contained in the server certificate from this site. Click [Add](#) to add a name in the Distinguished Names panel. See the following graphic:

Adding a Distinguished Name to a List



The dialog box titled "Add Distinguished Name" features a text input field labeled "Item" and two buttons at the bottom: "OK" with a green checkmark icon and "Cancel" with a red X icon.

If only the Common Name is required in the entry, type the common name as it occurs in the server certificate (for example, `www.example.com`). If other elements of the Distinguished name are required, then these need to be input

using standard X.509 syntax. The following example shows how a DN may be defined using this syntax:

- *.microsoft.com – the * at the beginning is a wildcard
- cn=www.example.com
- CN=*.example.com, OU=Research, O=Example Systems, Inc., C=US

The entries are case insensitive.

The following graphic shows examples of name entries that use the different input formats.

Examples of Distinguished Name Formats

Distinguished Names	
Item	
*.google.com	
cn=*.microsoft.com	
www.bbs.co.uk	

Common Names List

The Common Names lists display contains two panels with the lower panel displaying information that varies depending on the row selected in the upper panel. Each CN list occupies one row in the Common Names Lists pane. Tool icons on this panel in addition to the multi page tool allow the user to add or remove a list, or to clone a list. The remove and cloning icons are grayed out unless an entry in the table is selected. The difference between Common Names Lists and Distinguished Names Lists is that the Common Names Lists can only contain X.509 Common Names while the Distinguished Names Lists can contain any X.509 fields. Searching of Common Names Lists is optimized so that these lists can contain many thousands of entries. A typical use for Common Names Lists might be to prevent inspection of traffic to many different sites of a particular type – for example banking sites. Selecting the list in the upper panel causes the set of names in the list to be displayed in the lower panel.

Common Names Lists	
Name	
test	

Common Names	
Item	
*.google.com	

Maintaining large Common Names Lists using the WebUI is a time-intensive task. External tools that simplify and automate the management of such lists may be available to simplify this task.

IP Address Lists

The IP Addresses Lists display contains two panels with the lower panel displaying information that varies depending on the row selected in the upper panel. Each IP Addresses list occupies one row in the IP Addresses Lists panel. Searching of IP Address Lists is optimized so that these lists can contain many thousands of entries. A typical use for an IP Address List might be to prevent inspection of traffic to many different sites of a particular type based on the destination IP address of the hosts. Tools icons on this panel allow addition or removal of a list or cloning of a list. Selecting a list in the upper panel causes the set of addresses in the list to be displayed in the lower panel. IP addresses can be specified in three different formats:

- a. b. c. d – for example, 192. 168. 2. 10 (netmask of 255. 255. 255. 255 is implied)
- a. b. c. d/x – for example, 192. 168. 2. 1/24
- a. b. c. d: e. f. g. h – for example, 192. 168. 2. 1: 255. 255. 255. 224

Addresses are validated on input so the system will not allow input of an illegal IP address. Maintaining large IP Address Lists using the WebUI is a time-intensive task. External tools that simplify and automate the management of such lists may be available to simplify this task.

The following graphic shows the IP Addresses panel with three addresses entered, each using one of the three different input formats.

IP Addresses in Different Formats

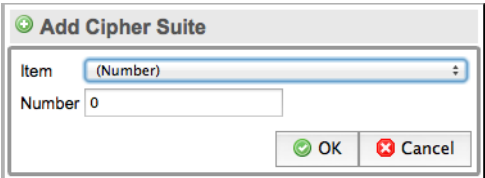
IP Addresses	
Item	
192.168.2.1/24	
10.1.1.54	
192.168.1.1:255.255.255.224	

Cipher Suites List

The Cipher Suites Lists display contains two panels with the lower panel displaying information that varies depending on the row selected in the upper panel. Each Cipher Suite list occupies one row in the Cipher Suites Lists panel. Tools icons on this panel allow addition or removal of a list or cloning of a list. Select a list in the upper panel to display the set of cipher suites in the list in the lower panel. When adding a cipher suite to a list, a dialog box appears that allows the cipher suite to be chosen from a drop down list or input as a number in decimal or hex format.

The following graphic shows the input box used to add a cipher suite.

Adding a Cipher Suite to a Cipher Suites List

A dialog box titled "Add Cipher Suite" with a green plus icon in the top left corner. It contains two input fields: "Item" with a dropdown menu showing "(Number)" and "Number" with a text box containing "0". At the bottom right are "OK" and "Cancel" buttons with green and red icons respectively.

The following graphics shows a list with three entries each using a different input format. The drop-down menu provides a list of all cipher suites using the name format: for example, TLS_RSA_WITH_DES_CBC_SHA.

Examples of Different Cipher Suite Formats

Cipher Suites	
Item	
TLS_RSA_WITH_DES_CBC_SHA	
47	
0x2E	

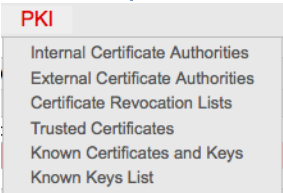
PKI Management

The PKI menu contains six options that manage certificates, keys and the creation of lists of certificates and keys. Each of the menu options is described below.

WARNING! A user must have the Manage PKI role to make changes to the certificates and keys on the system. Some features of the PKI menu will not be available to users without the Manage PKI role.

The following graphic shows the PKI menus options.

PKI Menu Options

A screenshot of the PKI menu. The menu is titled "PKI" in red. It lists six options: "Internal Certificate Authorities", "External Certificate Authorities", "Certificate Revocation Lists", "Trusted Certificates", "Known Certificates and Keys", and "Known Keys List".

Internal Certificate Authorities

The Internal Certificate Authorities screen has a single panel where Certificate Authorities for Certificate Re-sign decryption are created, imported, exported and managed. The panel has the following tool icons: multi-page icon, generate certificate, install certificate, export certificate, view certificate details, add certificate, delete certificate, edit, and refresh.

The different ways an Internal CA can be added are described in [Installing a CA for Certificate Re-Sign](#) on page 83. Multiple internal Certificate Authorities can be configured and stored in the system. A segment, ruleset, or definition can determine which internal CA is used to re-sign a server certificate when an SSL session is being decrypted using certificate re-sign. The choice of internal CA can be configured to depend on details of the server certificate; different internal CAs can be used for traffic going to different servers over the same segment.

External Certificate Authorities

The External Certificate Authorities Lists displays two panels; the lower panel displays information depending on the row selected in the upper panel. Each External Certificate Authorities list occupies one row in the External Certificate Authorities Lists panel. Tools icons on this panel allow addition or removal of a list or cloning of a list. Selecting a list in the upper panel causes the set of External CA certificates in the list to be displayed in the lower panel.

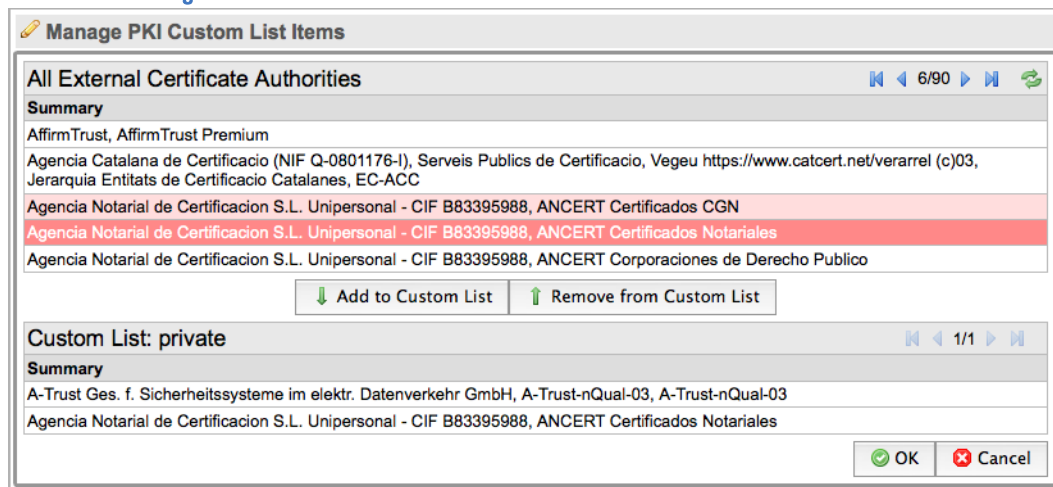
The system has a default list installed, the [all-external-certificate-authorities](#) list. This contains the set of publicly trusted CA certificates that are distributed with Internet Explorer and Firefox browsers. Select this list in the upper panel to display the details of the CA certificates in the lower panel.

The External Certificate Authorities panel has the following tool icons: multi-page icon, view certificate details, add certificate, delete certificate, and refresh. This allows additional trusted CA certificates to be added to the list or for existing CA certificates in the system to be deleted.

Use the add button on the External Certificate Authorities Lists panel to create and add a custom list. After this list is created, it can be selected and CA certificates from the [all-external-certificate-authorities](#) list can be copied to the custom list. The custom list is always a subset of the [all-external-certificate-authorities](#) list and cannot contain entries that are not present in the [all-external-certificate-authorities](#) list. When a custom list is selected and the add button in the lower panel is clicked, a dialog box appears allowing keys in the default list to be added to the custom list.

The following graphic shows an example where two CA certificates from the `all-external-certificate-authorities` list have been added to a custom list called `private`.

Creating a Custom External Certificate Authorities List



One of the entries that has been included in the `private` list is a private CA certificate that had previously been imported to the `all-external-certificate-authorities` list – the Example Systems CA.

The clone feature on the External Certificate Authorities Lists panel is used to clone an existing list and save it with a new name. It is often quicker to clone an existing custom list and then add or remove certificates to the new version produced by the clone tool.

Certificate Revocation Lists

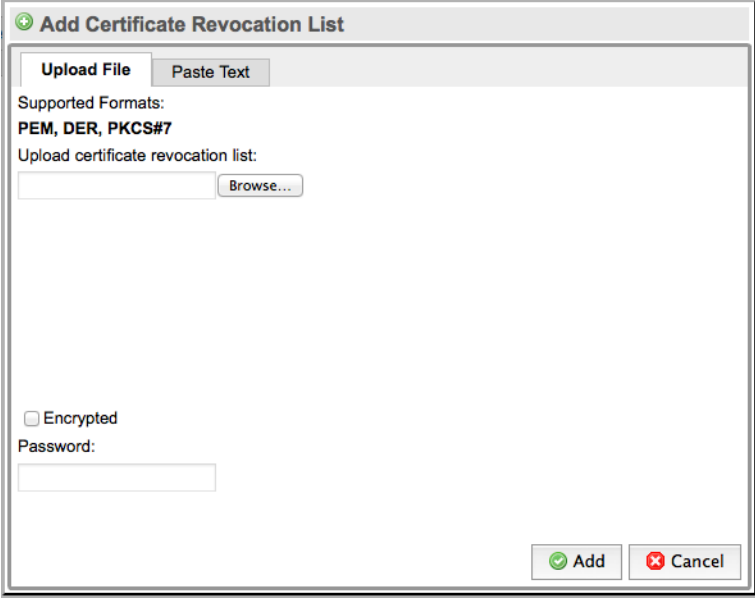
The Certificate Revocation Lists displays two panels; the lower panel displays information depending on the row selected in the upper panel. Each Certificate Revocation List occupies one row in the List of Certificate Revocation Lists panel. Tools icons on this panel allow addition or removal of a list or cloning of a list. Selecting a list in the upper panel causes the set of CRLs in the list to be displayed in the lower panel.

The system has a default list installed, the `all-certificate-revocation-lists` list. This list is initially empty. Select this list in the upper panel to display details of the CRLs in the lower Certificate Revocation Lists panel. Select this list and click on the add button in the lower Certificate Revocation Lists panel to open up a dialog box that allows import of a CRL.

The Certificate Revocation Lists panel has the following tool icons: multi-page icon, view CRL details, add CRL, delete CRL, and refresh. This allows CRLs to be

added to the list or for existing CRLs in the system to be deleted. The following graphic shows the import CRL dialog box.

Add Certificate Revocation List Box

The image shows a web-based dialog box titled "Add Certificate Revocation List". It has two tabs: "Upload File" (selected) and "Paste Text". Under the "Upload File" tab, it lists "Supported Formats: PEM, DER, PKCS#7". Below this, it says "Upload certificate revocation list:" followed by a text input field and a "Browse..." button. At the bottom left, there is a checkbox labeled "Encrypted" and a "Password:" label with a corresponding text input field. At the bottom right, there are two buttons: "Add" (with a green checkmark icon) and "Cancel" (with a red X icon).

If the CRL file being imported is encrypted and protected with a password, then the password will need to be entered in the Password field on the box.

The **Add** button on the List of Certificate Revocation Lists panel can be used to create and add a custom list. After this list is created, it can be selected and CRLs from the [all -certificate-revocation-lists](#) list can be copied to the custom list. The custom list is always a subset of the [all -certificate-revocation-lists](#) list and cannot contain entries that are not present in the [all -certificate-revocation-lists](#) list. Select a custom list and click the **Add** button in the lower panel to display a dialog keys in the default list that can be added to the custom list.

The clone feature on the List of Certificate Revocation Lists panel is used to clone an existing list and save it with a new name. It is often quicker to clone and existing custom list and then add or remove CRLs to the new version produced by the clone tool.

Trusted Certificates

The Trusted Certificates Lists displays two panels; the lower panel displays information depending on the row selected in the upper panel. The Trusted Certificates List occupies one row in the Trusted Certificates Lists panel. Tools icons on this panel allow the addition or removal of a list or the cloning of a list. Selecting a list in the upper panel causes the set of certificates in the list to be displayed in the lower panel.

The system has a default list installed, the `all-trusted-certificates` list. This list is initially empty. Select this list in the upper panel to display details of the certificates in the lower Trusted Certificates panel. Select this list and click on the add button in the lower Trusted Certificates panel to import a certificate.

The Trusted Certificates panel has the following tool icons: multi-page icon, view certificate details, add certificate, delete certificate, and refresh. This allows certificates to be added to the list or existing certificates in the system to be deleted.

Click the **Add** button on the Trusted Certificates Lists panel to create or add a custom list. After this list is created, copy certificates from the `all-trusted-certificates` list to the custom list. The custom list is always a subset of the `all-trusted-certificates` list and cannot contain entries that are not present in the `all-trusted-certificates` list. When a custom list is selected and the **Add** button in the lower panel is clicked, a dialog box appears allowing keys in the default list to be added to the custom list.

The clone feature on the External Certificate Authorities Lists panel is used to clone an existing list and save it with a new name. It is often quicker to clone an existing custom list and then add or remove certificates to the new version produced by the clone tool.

Known Certificates and Keys

The Known Certificates and Keys list displays two panels; the lower panel displays information depending on the row selected in the upper panel. Each Known Certificates and Keys List occupies one row in the Known Certificates and Keys Lists panel. Tools icons on this panel allow addition or removal of a list or cloning of a list. Select a list in the upper panel to display the set of certificates with keys in the list in the lower panel.

The system has a default list installed: the `all-known-certificates-with-keys` list. This list is initially empty. Select this list in the upper panel to display the details of the certificates with keys in the lower Known Certificates with Keys panel. Select this list and then click on the **Add** button in the lower Known Certificates with Keys panel to open up a dialog box that allows import of a certificate into the default list.

The Known Certificates with Keys panel has the following tool icons: multi-page icon, view certificate details, add certificate, delete certificate, and refresh. This allows certificates and keys to be added to the list or for existing certificates and keys in the system to be deleted.

Click the **Add** button on the Trusted Certificates Lists panel to create or add a custom list. After this list is created, copy certificates from the `all-known-certificates-with-keys` list to the custom list. The custom list is always a subset of the `all-known-certificates-with-keys` list and cannot contain entries that are not present in the `all-known-certificates-with-keys` list. When a custom list is selected and the **Add** button in the lower panel is

clicked, a dialog box appears allowing keys in the default list to be added to the custom list.

The clone feature on the Known Certificates with Keys Lists panel can be used to clone an existing list and save it with a new name. It is often quicker to clone an existing custom list and then add or remove certificates to the new version produced by the clone tool.

Known Keys List

The Known Keys list displays two panels; the lower panel displays information depending on the row selected in the upper panel. Each Known Keys List occupies one row in the Known Keys Lists panel. Tools icons on this panel allow addition or removal of a list or cloning of a list. Select a list in the upper panel to display the set of certificates with keys in the list in the lower panel.

The system has a default list installed: the `al l -known-keys` list. This list is initially empty. Select this list in the upper panel to display the details of the keys in the list in the lower Known Certificates with Keys panel. Select this list and then click on the **Add** button in the lower Known Keys panel to open up a dialog box that allows import of a key into the default list.

The Known Keys panel has the following tool icons: multi-page icon, view key details, add key, delete key, and refresh. This allows keys to be added to the list or for existing keys in the system to be deleted.

Click the **Add** button on the Trusted Certificates Lists panel to create or add a custom list. After this list is created, copy certificates from the `al l -known-keys` list to the custom list. The custom list is always a subset of the `al l -known-keys` list and cannot contain entries that are not present in the `al l -known-keys` list. When a custom list is selected and the **Add** button in the lower panel is clicked, a dialog box appears allowing keys in the default list to be added to the custom list.

The clone feature on the Known Keys Lists panel can be used to clone an existing list and save it with a new name. It is often quicker to clone an existing custom list and then add or remove keys to the new version produced by the clone tool.

Platform Management

The Platform Management menu contains ten options. These are all described in the sections that follow. This menu includes tools to view and manage the platform and to configure and manage access to the platform network management features. Platform management also includes managing user accounts and performing updates to the system software.

Information

The Information screen initially displays two panels and a button to access additional information. The two panels have only refresh buttons because they provide visibility of data but no ability to enter or change data.

The following graphic shows the Software Versions panel which provides details of the software versions of the various software modules within the system.

Platform Information -Software Versions

Software Versions	
SSL Appliance Linux Distribution:	3.5.1-647
SSL Appliance Linux Distribution (WebUI VM):	3.5.0-195
Linux Kernel:	2.6.38-8-server
Netronome Flow Processor Drivers:	nfp-bsp-release-2011.11.1-20111107233052.ns
GERD Software:	gerchr-src-1.0.0-r144
Netronome Flow Driver:	2.6.0-2195
Netronome Standard Library:	1.0.3-512
Netronome Security Module:	1.3.0-1500
SSL Appliance Software:	1.6.0-2362
SSL Appliance WebUI:	1.1.0-238

The SSL Appliance Linux Distribution value, in this example, **3.5.1-647**, is the most important element here; this is the version number of the software release that is running on the system. Sourcefire personnel may request the details from this panel when providing support for the device. The following graphic Shows the Chassis FRU Info panel.

Platform Information -Chassis Data

Chassis FRU Info	
Chassis Part Number	Chassis Serial Number
NFPP-1U-AC	515-11012200400025

Sourcefire personnel may request the details from this panel when providing support for the device.

Click the **Show Advanced** button to display an additional set of view-only panels that provide data on different hardware elements of the system. Sourcefire personnel may request the details from these panels when providing support for the device. Panels provide details for the following hardware components of the system:

- Midplane VPD Info: midplane that connects NetMods to switch and switch to NFE card
- Switch Board VPD info: switch that plugs into midplane
- NetMod VPD Info: details on the NetMods plugged in to the system
- CPU Info: details on the CPUs installed on the system motherboard
- NFE VPD Info: details on the NFE cards installed in the system

Management Network

The Management Network screen has a single panel that allows configuration of the management network settings. The panel has edit and refresh tool icons.

The system can use either a fixed IP address or acquire an IP address using DHCP. For DHCP to work the management ethernet must be connected to a network with a working DHCP server.

The following graphic shows the panel containing data for a system that is configured to use a static IP address and which currently still has the default Hostname.

Management Network Panel

Management Network	
MAC Address:	00:15:4d:00:85:c5
MTU:	1500
DHCP:	True
IP Address:	192.168.3.214
Netmask:	255.255.255.0
Default Gateway:	192.168.3.1
Hostname:	localhost.localdomain
Primary Nameserver:	
Secondary Nameserver:	

The following graphic shows the configuration dialog box used to adjust the network settings.

Network Management Configuration Box

Edit Management Network	
MTU	<input type="text" value="1500"/>
DHCP	<input checked="" type="checkbox"/>
IP Address	<input type="text" value="192.168.3.214"/>
Netmask	<input type="text" value="255.255.255.0"/>
Default Gateway	<input type="text" value="192.168.3.1"/>
Hostname	<input type="text" value="localhost.localdomain"/>
Primary Nameserver	<input type="text"/>
Secondary Nameserver	<input type="text"/>
<input type="button" value="OK"/> <input type="button" value="Cancel"/>	

When you check the DHCP check box, the fields for IP Address, Netmask and Default Gateway are grayed out. [Configuring Management Network Settings](#) on page 79 includes more details on configuring the management network settings.

Remote Logging

Use remote logging to send log messages to a remote syslog server. To enable the system to send log messages to a remote syslog server as well as to the internal log file you must enable remote logging by checking the checkbox, provide the IP address and port number, select the protocol from the dropdown menu, and select whether to send all logs or only warnings and errors. You can configure up to two remote Syslog servers.

Date/Time

The Date/Time screen has a single panel that allows configuration of the system time and date settings. The panel has edit and refresh tool icons. In addition to setting the time and date you can configure the time zone and set NTP to synchronize the system to a network time server.

The following graphic shows the panel for a system that is configured to use NTP and is located in the UK time zone.

Date/Time Panel

Date/Time	
Date:	2011-12-8
Time:	15:44:22
Timezone:	UTC
NTP Enabled:	False
Primary NTP:	
Secondary NTP:	

Click on the edit tool to open up a dialog box that allows the settings to be changed. Reboot the system to apply changes to the date/time settings. More details on setting the date/time can be found in [Configuring System Date/Time and Time Zone](#) on page 78.

Users

The Users menu has a single panel with tool icons for multi-page, add, edit, delete and refresh. Only users with Manage Appliance or Manage PKI roles can make changes to the user accounts on the system.

The following graphic show the Users panel for a system that has three user accounts configured.


Managing User Accounts on the System

User Management		
User ID	Full Name	Roles
admin	admin	Manage Policy, Manage Appliance, Auditor, Manage PKI

Each account has a different set of roles associated with it. More details on creating user accounts and on the meaning of different roles can be found in [Configuring Management Users](#) on page 80.

TACACS Servers

A Cisco ACS system using TACACS+ can be used to remotely authenticate access to the SSL appliance management webUI. This menu option allows the system to be configured to use TACACS+ to communicate with a Cisco ACS.

Below is the TACACS server panel with an entry already present, initially the table is empty. Use the  button to create an entry.

TACACS Servers				
IP	Port	Retries	Network Timeout	Retry Timeout
172.18.0.201	49	0	0	0

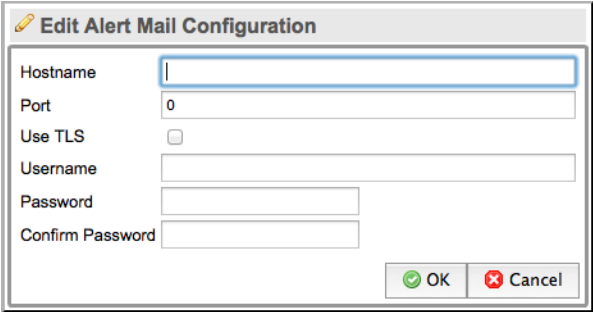
When creating the entry you need to provide the details shown. The secret value needs to match the secret value configured on the ACS server. If TACACS is in use, the login box on the webUI will change to include a drop down menu that lets the user choose if they want to be authenticated remotely or locally as shown above.

Alerts

The Alerts menu contains two panels that allow configuration of the email details that the system will use to send out alerts and of the events that are to be monitored and the conditions under which an alert is generated.

The upper Alert Email Configuration panel includes an edit tool icon and is used to configure details of the email system that will be used to send out alerts. Click [Edit](#) to display a dialog box as shown in the following graphic.

Email Configuration for Alert System



The dialog box titled "Edit Alert Mail Configuration" contains the following fields and controls:

- Hostname:** A text input field.
- Port:** A text input field with the value "0".
- Use TLS:** A checkbox, currently unchecked.
- Username:** A text input field.
- Password:** A text input field.
- Confirm Password:** A text input field.
- Buttons:** "OK" (with a green checkmark icon) and "Cancel" (with a red X icon).

To configure an email alert, provide the following items:

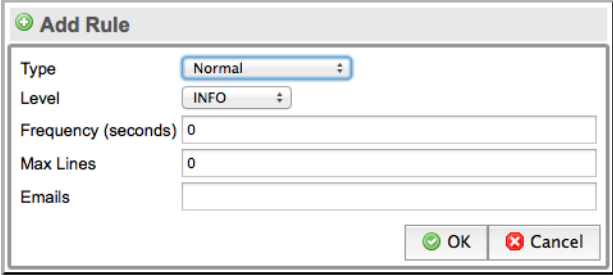
- **Hostname:** this is the name or IP address of the SMTP server used to send email
- **Port:** this is the port number on the SMTP server that is used to send email
- **Use TLS:** check box to enable or disable the use of encryption (TLS) when sending email

- Username: the username of the account being used to send email
- Password: the password for the account being used to send email

WARNING! If your enterprise is using Google Apps for email then the correct SMTP Server Address is `aspmx.l.google.com`, not `smtp.gmail.com`. Alerts can only be sent to users on the same domain using this SMTP configuration.

Configure individual alerts in the lower panel. Each alert can be triggered by a specific set of conditions and can be sent to one or more email recipients. Click **Add** in the lower panel to opens up a dialog box and configure the alert.

Add Alert to the System



The Alert type can be:

- **Harddrive Full**: generated if out of disk space
- **Normal**: generated if conditions specified in alert are met
- **Periodic**: generated at regular time intervals
- **Unclean Shutdown**: generated if last system shutdown was not clean

The Level type can be:

- **ERROR**
- **FATAL**
- **INFO**
- **WARNING**

These levels correspond to entries in the system log files. For example, if the Level is set to **FATAL**, an alert will be generated when a message with a **FATAL** level is added to the system log.

The Frequency box controls how frequently the alert message are sent and the Max Lines box controls how many lines from the system log are included in the email that is sent.

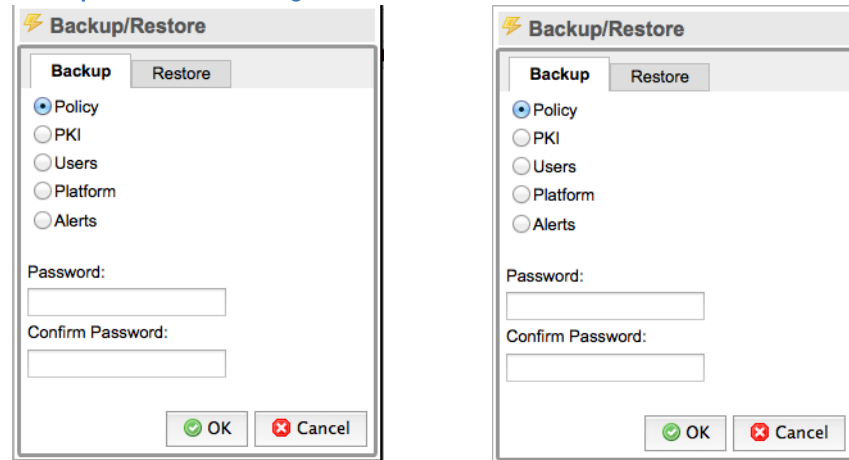
Specify which email addresses will receive alerts in the **Emails** box.

Backup/Restore

Save or restore elements of the system configuration from a remote storage system in the Backup/Restore menu.

The following graphics show the backup dialog box and the restore dialog box.

Backup and Restore Dialog Boxes

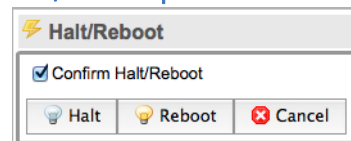


Select the item to be backed up or restored by clicking the radio button for that item. You must provide a password backing up and restoring data.

Halt/Reboot

Halt or reboot the system in the Halt/Reboot menu. The following graphic shows the dialog box.

Halt/Reboot Option



The Halt and Reboot buttons are grayed out until the confirm check box is checked.

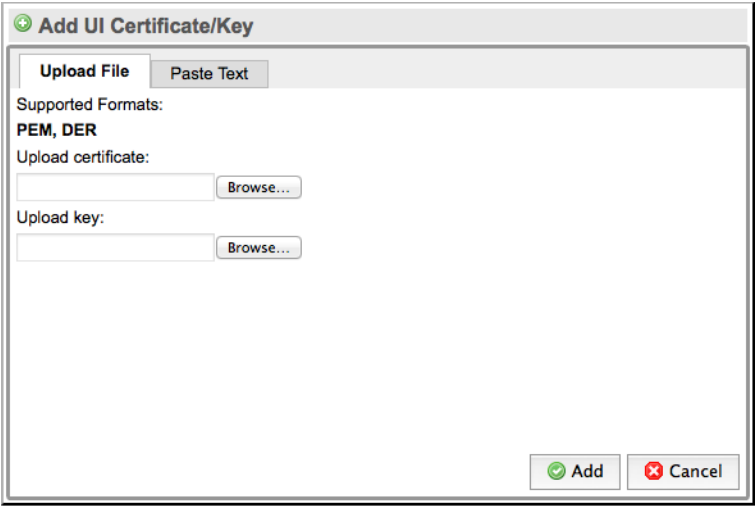
WARNING! If the system is halted, it will require physical presence to power it on from the front panel power switch.

Import UI Certificate/Key

This menu option allows a signed SSL server certificate to be imported for use by the webserver that provides the webUI management for the system. By default the system uses a self-signed server certificate which will cause warnings from

browsers; see [Browser Configuration](#) on page 102 for details. The following graphic shows the dialog box used to import a certificate for use by the WebUI.

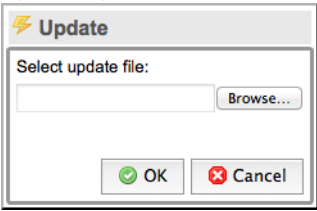
Import Certificate for WebUI

A dialog box titled "Add UI Certificate/Key" with a green plus icon. It has two tabs: "Upload File" (selected) and "Paste Text". Under "Supported Formats:", it lists "PEM, DER". There are two sections: "Upload certificate:" with a text input field and a "Browse..." button, and "Upload key:" with a text input field and a "Browse..." button. At the bottom right are "Add" (with a green checkmark) and "Cancel" (with a red X) buttons.

Update

The update menu option is used to load and apply an update file that will update the system software. Update files are digitally signed and are checked before they are applied to the system. Invalid update files will not be applied. The following graphic shows the update dialog box.

Update System Box

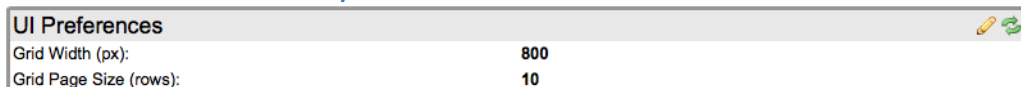
A dialog box titled "Update" with a yellow lightning bolt icon. It contains a label "Select update file:" above a text input field and a "Browse..." button. At the bottom are "OK" (with a green checkmark) and "Cancel" (with a red X) buttons.

The [Choose File](#) button opens a window that allows the user to browse their system and to select the update file that is to be used. After the [OK](#) button is pressed, the file is checked and, if valid, will be copied to the system and then applied.

Preferences

Configure the UI screen layout in the Preferences menu. The following graphic shows the panel with the default values showing for the grid width and number of rows.

Preference for WebUI Layout

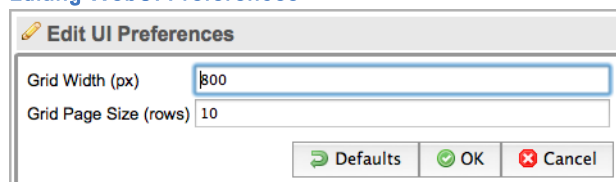


The image shows a 'UI Preferences' panel. It has a title bar with a pencil icon and a refresh icon. Below the title bar, there are two rows of settings. The first row is 'Grid Width (px):' with a value of '800'. The second row is 'Grid Page Size (rows):' with a value of '10'.

UI Preferences	
Grid Width (px):	800
Grid Page Size (rows):	10

Click **Edit** to change the grid width and number of rows, or force them back to the system defaults. This is shown in the following graphic.

Editing WebUI Preferences



The image shows an 'Edit UI Preferences' dialog box. It has a title bar with a pencil icon. Below the title bar, there are two input fields. The first is 'Grid Width (px)' with a value of '800'. The second is 'Grid Page Size (rows)' with a value of '10'. At the bottom right, there are three buttons: 'Defaults' (with a circular arrow icon), 'OK' (with a green checkmark icon), and 'Cancel' (with a red X icon).

Edit UI Preferences	
Grid Width (px)	800
Grid Page Size (rows)	10


Defaults OK Cancel

WARNING! Multi-page panels have a built in multiplier that is used in conjunction with the number of rows value that is configured as the default. For example, the SSL Statistics panel has a multiplier of 1.6. With the default row setting of 10, this will mean there are 16 rows displayed in the SSL statistics panel. If the default row count was set to 20 then the SSL Statistics panel would have 32 rows.

User Management

A user can change their password or log out in the User menu.

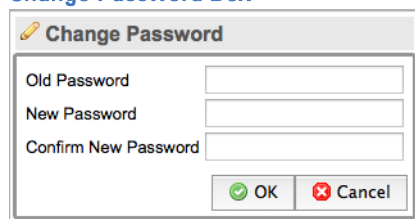
User Menu



Change Password

The following graphic shows the change password dialog box.

Change Password Box

A screenshot of the 'Change Password' dialog box. The title bar says 'Change Password' with a yellow key icon. Inside the dialog, there are three text input fields labeled 'Old Password', 'New Password', and 'Confirm New Password'. At the bottom right, there are two buttons: 'OK' with a green checkmark icon and 'Cancel' with a red X icon.

Input your current password and then the new password to change your password. Passwords must be at least 8 characters long and must contain at least one alpha character and one numeric character.

Logout

Select the logout option to log off.

Chapter 6

Troubleshooting the System

IMPORTANT! Please read through all the information in this section of the document before contacting Sourcefire support.

Supported Network Protocols and Frame Encapsulations

The SSL appliance supports SSL processing on TCP in IPv4 and IPv6. The IP packet must be encapsulated in an Ethernet-II frame, with an optional VLAN tag (802.1Q or 802.1ad).

Network traffic for all other protocols and frame encapsulations are not sent to the SSL processing engine, including the following:

- QinQ (VLAN double tagging)
- Cisco ISL
- MPLS
- GRE
- IP-in-IP
- UDP
- ICMP
- ARP
- SOCKS

- DSSL
- IPsec

Supported SSL/TLS Versions

This version of the SSL appliance only supports SSL 3.0, TLS 1.0, TLS 1.1, and TLS 1.2. There is no support for SSL 2.0. If SSL 2.0 traffic is encountered, the SSL appliance will either Cut Through or Reject the flow according to the [Undecryptable SSL Handling](#) parameter in the SSL Inspection Policy. Note that SSL 2.0 ClientHello messages are supported as long as the rest of the SSL handshake is completed using version 3.0 or above (more detail on this compatibility mode can be found in Section E.1 of RFC4346).

Support for Client Certificates

The SSL appliance supports decrypting SSL sessions with client certificates only if the action in the inspection policy is [Decrypt: server key is known](#) and RSA is used as the key exchange algorithm. The CertificateVerify SSL handshake message sent after the client certificate is digitally signed by a key known only to the client, because the CertificateVerify message cannot be modified and no part of the SSL handshake can be modified.

SSL sessions using client certificates and the RSA key exchange in known server key mode are decrypted as usual. The SSL appliance rejects all other sessions with client certificates unless the session uses a cipher suite which is not in the list of [Supported Cipher Suites](#) on page 143. SSL sessions rejected because of a client certificate appear in the SSL session log with an [Error](#) event value and [Reject](#) action. The SSL appliance dataplane log also shows an error message with details about the rejected session.

To prevent sessions with client certificates from being rejected, the Inspection Policy must have a rule that will cut through the specific session based on a combination of common name, destination IP address/mask, and destination TCP port. Alternatively, you can cut the sessions through using the Traffic Diversion Policy.

Supported Cipher Suites

The Supported Cipher Suites list contains all the cipher suites that are supported by the SSL appliance and shows which can be inspected when inline and when in

passive-tap mode. Any cipher suites that are not supported will be handled by the policies configured for undecryptable traffic.

Supported Cipher Suites

Cipher Suite	Inline	Passive-Tap	ID
TLS_NULL_WITH_NULL_NULL	Yes	Yes	0x0000
TLS_RSA_WITH_NULL_MD5	Yes	Yes	0x0001
TLS_RSA_WITH_NULL_SHA	Yes	Yes	0x0002
TLS_RSA_WITH_RC4_128_MD5	Yes	Yes	0x0004
TLS_RSA_WITH_RC4_128_SHA	Yes	Yes	0x0005
TLS_RSA_WITH_DES_CBC_SHA	Yes	Yes	0x0009
TLS_RSA_WITH_3DES_EDE_CBC_SHA	Yes	Yes	0x000A
TLS_DHE_RSA_WITH_DES_CBC_SHA	Yes	No	0x0015
TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	Yes	No	0x0016
TLS_DH_Annon_WITH_DES_CBC_SHA	Yes	No	0x001A
TLS_DH_Annon_WITH_3DES_EDE_CBC_SHA	Yes	No	0x001B
TLS_RSA_WITH_AES_128_CBC_SHA	Yes	Yes	0x002F
TLS_DHE_RSA_WITH_AES_128_CBC_SHA	Yes	No	0x0033
TLS_DH_Annon_WITH_AES_128_CBC_SHA	Yes	No	0x0034
TLS_RSA_WITH_AES_256_CBC_SHA	Yes	Yes	0x0035
TLS_DHE_RSA_WITH_AES_256_CBC_SHA	Yes	No	0x0039
TLS_DH_Annon_WITH_AES_256_CBC_SHA	Yes	No	0x003A
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA	Yes	Yes	0x0041
TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA	Yes	No	0x0045
TLS_DH_Annon_WITH_CAMELLIA_128_CBC_SHA	Yes	No	0x0046

Supported Cipher Suites (Continued)

Cipher Suite	Inline	Passive-Tap	ID
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	Yes	Yes	0x0084
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA	Yes	No	0x0088
TLS_DH_Annon_WITH_CAMELLIA_256_CBC_SHA	Yes	No	0x0089
TLS_ECDHE_RSA_WITH_NULL_SHA	Yes	No	0xC010
TLS_ECDHE_RSA_WITH_RC4_128_SHA	Yes	No	0xC011
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	Yes	No	0xC012
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	Yes	No	0xC013
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	Yes	No	0xC014
TLS_ECDH_Annon_WITH_NULL_SHA	Yes	No	0xC015
TLS_ECDH_Annon_WITH_RC4_128_SHA	Yes	No	0xC016
TLS_ECDH_Annon_WITH_3DES_EDE_CBC_SHA	Yes	No	0xC017
TLS_ECDH_Annon_WITH_AES_128_CBC_SHA	Yes	No	0xC018
TLS_ECDH_Annon_WITH_AES_256_CBC_SHA	Yes	No	0xC019
SSL_RSA_FIPS_WITH_DES_CBC_SHA	Yes	Yes	0xFEFE
SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA	Yes	Yes	0xFEFF
SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA	Yes	Yes	0xFFE0
SSL_RSA_FIPS_WITH_DES_CBC_SHA	Yes	Yes	0xFFE1

There is no support for the outdated export version of the cipher suites. There is also no support for SHA256 MAC, static DH (Diffie-Hellman) key exchange, DSS (Digital Signature Standard) authentication, or ECC (Elliptic Curve) key exchange.

IMPORTANT! When operating in passive-tap mode, there are some cipher suites that cannot be inspected (for example, Ephemeral and Anonymous DH key exchanges). When operating in inline modes, it is possible to inspect SSL sessions using Ephemeral and Anonymous DH key exchanges.

SSL sessions using unsupported cipher suites appear in the SSL session log with an **Undecryptable** event value. The action taken depends on the **Undecryptable SSL Handling** policy option and is either **Cut through**, **Drop**, or **Reject**.

Note that there are no restrictions on cipher suites for policies with actions that do not involve inspecting the traffic. For example, you can have a policy that prevents SSL traffic using ECC from setting up connections across the network.

Support for SSL Record Layer Compression

The SSL specification allows the SSL record layer compression to use an algorithm negotiated through the ClientHello and ServerHello handshake messages. The current version of the SSL appliance does not support SSL record layer compression, and all such SSL sessions will be marked as **Undecryptable** in the SSL session log. The action taken on these sessions is determined by the Undecryptable SSL Handling policy option.

Support for Stateless Session Resumption (RFC5077)

The SSL appliance supports stateless session resumption as outlined in RFC5077. Stateless sessions are typically used by content providers that balance high loads between multiple servers. An example of this is Google Mail (www.gmail.com).

Steps to Troubleshoot SSL Decryption

If none of the incoming SSL sessions are decrypted, follow the steps outlined below.

Monitor Network Port Statistics

Verify that network traffic is received on the network ports of the SSL appliance being used by the active segment. The Monitor/Dashboard screen on the WebUI provides the required information in the Segment Status and Network Interfaces panels.

Monitor the SSL Statistics

Verify that SSL sessions reach the SSL processing engine of the SSL appliance. The SSL Statistics option on the Monitor WebUI menu provides the required information. If you can see the counts for detected SSL session increasing, SSL traffic is being detected by the system.

Monitor the SSL Session Log

Verify that SSL sessions are recorded in the SSL session log and have the correct status. The [SSL Session Log](#) option under the [Monitor](#) menu provides the required information. First, ensure that the SSL Session Log is enabled for the segment being used. Next, confirm that the SSL sessions appear in the session log: ensure that you are viewing the first page of session log data and press the [Refresh](#) button and you should see new entries appear at the top of the page. Appropriate values in the [Action Taken](#) column confirm that the SSL sessions are being decrypted. The session log identifies the segment and entry with the segment you are troubleshooting. This can be found on the Policies / Segment screen.

Verify Inspection Policy Configuration

Verify that the rules specified in the ruleset being used on the segment of interest are set up to inspect the traffic that you are interested in. See [Rulesets](#) on page 116 for more details.

Known Server Versus Trusted Server Certificates

The server's private key and certificate must be loaded into the Known Certificates and Keys store before inspecting traffic to that server. Known Server Certificates are implicitly trusted and need not be signed by a CA trusted by the SSL appliance.

Do **not** install server certificates in the Trusted Certificates store if you have the private key for that server. Instead store those certificates in the Known Certificates and Keys store. The Trusted Certificates store is used only to solve specific certificate validation problems, specifically, trusting self-signed certificates or trusting certificates for which you don't want to install the CA certificate chain. See [PKI Management](#) on page 126.

Caveats when Enabling/Disabling SSL Inspection

Immediately after you connect a segment to the network or activate inspection, you may not be able to decrypt some SSL flows. Such flows appear in the SSL session log (if activated) with a [Cut through](#) action and an [Uncached](#) certificate

common name (CN), and are handled according to the [Uncached SSL Session Handling](#) policy option. This happens because the flows are reusing an SSL session established before the SSL appliance was put inline, so the SSL appliance did not see the original full handshake and does not have the SSL session state cached.

An SSL session is established using a full SSL handshake, during which the peers negotiate the cryptographic state necessary to encrypt and decrypt traffic. SSL clients, such as web browsers and email clients, cache the cryptographic state and may reuse the session multiple times in later SSL flows. Similarly, the SSL appliance inspects the full handshake, caches the session state, and uses it to inspect flows reusing the same session. If the full handshake occurred before the appliance was put inline, it cannot decrypt flows reusing that session. Most servers allow sessions to be reused only for a few hours, after which they force clients to establish new sessions. Therefore, the SSL session log may show [Uncached](#) sessions for a few hours after installing the device on the network or activating inspection. After the client and server establish a new SSL session, the SSL appliance can decrypt that session and all subsequent sessions between the same client and server.

Note that SSL clients may report SSL session failures if you disconnect the SSL appliance. If an application (for example, Microsoft Outlook) supports SSL session reuse, it will report a failure when it tries to reuse the SSL session. The SSL session fails because the full SSL handshake was used to establish the initial SSL session, and the SSL appliance was acting as a man-in-the-middle (MITM). The session that the client is trying to reuse was actually a session from the client to the appliance rather than to the server. The client does not know this because the SSL appliance is a transparent MITM. However, if the MITM is removed and the client attempts session reuse, the request goes to the server and the server cannot reuse this session because it does not recognize it.

Generating the Internal CA Certificates

Inspecting SSL sessions in any of the inline modes requires at least one internal CA certificate and private key, unless only Known Key decryption is used. The SSL appliance can generate the internal CA private key and either a self-signed certificate or a Certificate Signing Request (CSR) that can be forwarded to another CA. If using the CSR option, it is important to note that public CA companies such as Verisign are unlikely to issue intermediate CA certificates for use in the SSL appliance. See [Installing a CA for Certificate Re-Sign](#) on page 83 and [Internal Certificate Authorities](#) on page 126 for more details.

Access to Microsoft Windows Update Denied

When trying to access the Microsoft Windows update service through the SSL appliance, an error message may be displayed by Internet Explorer, and the update service will fail.

This is because the CA of the certificate presented by the update website server does not belong to a Microsoft server. The update is aborted with an error. To allow the updates to continue, add an SSL Inspection Policy for the certificate Common-Name *.update.microsoft.com with an action of **Cut Through** without decrypting.

IMPORTANT! Default rules are created by the system to deal with this and other sites where attempting to inspect SSL traffic will cause a problem.

Issues with Alerts

If you fail to receive email alerts, check the system log file for errors. The following may also prevent email from being sent or delivered:

- Check that, if your SMTP server requires authentication, the username and password specified in the SMTP Server Settings section is correct.
- Check that you are using the correct port for the specified SMTP server. Some servers are configured not to use the default port 25.
- Ensure that the SSL appliance has a fully qualified domain name (FQDN). Some SMTP servers require that the sender have a FQDN.
- Ensure that all email addresses are correct.
- Use `aspmx.l.google.com` instead of `smtp.gmail.com` as the SMTP server for Google Apps. Note that alerts can only be sent to users on the same domain with this SMTP configuration.

Procedure for Reporting an Issue

The first step in reporting an issue is to capture diagnostics using the WebUI. See [Diagnostics](#) on page 113 for details on how to generate diagnostic files.

Support engineers may request further diagnostic information such as SSL statistics, non-SSL statistics, and the SSL session log (if enabled). The engineers cannot request a copy of the PKI store because it may contain sensitive key material.

Preparing for Hardware Diagnostics or Maintenance

Support engineers may request advanced hardware diagnostics, or ask that certain firmware be upgraded. Before any upgrade, the SSL appliance must be put into a state where no traffic reaches the internal network interface, and packet processing engines are disabled. If this is required, appropriate directions will be given by the support engineer.

Chapter 7

Safety Information

Be sure to read the separate Safety Notice included in the SSL appliance packaging.

Safety Instructions

Please read all of the following instructions regarding the Sourcefire SSL appliance carefully.

- Ventilation: The Sourcefire SSL appliance vents (on the front panel) and the fan openings on the back panel are provided for ventilation and reliable operation of the product and to protect it from overheating. These openings must not be blocked or covered. This product must not be placed in a built-in installation unless proper ventilation is provided.
- Power Cords

WARNING! The power supply cords are used as the main disconnect device. Ensure that the socket outlet is located or installed near the equipment and is easily accessible. The SSL appliance has dual redundant power supplies that are powered by two separate power cords. Always disconnect BOTH cords to remove power from the unit.

WARNING! Do not disassemble this product. Return the SSL appliance to Sourcefire for service or repair work. Opening or removing covers may expose the user to dangerous voltage or other risks. Incorrect assembly can cause electric shock when this appliance is subsequently used.

IMPORTANT! Opening the cover voids the warranty.

Rack Mounting the Equipment

If the SSL appliance is to be installed in an equipment rack, please follow these precautions:

- Keep the ambient temperature around the appliance (which may be higher than the room temperature) within the operational limits specified in [Product Specifications](#) on page 13.
- Provide the unit with sufficient airflow.
- Do not overload the electrical circuits; consider the nameplate ratings of all the connected equipment and provide sufficient over current protection.
- Ground the equipment properly.
- Never place any objects on top of the appliance.

Chapter 8

Licenses and Licensing Terms

The following sections include details of license terms and conditions that apply to the Netronome SSL Inspector as well as licenses details for third party software that is included in the product.

GNU General Public License

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc. 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the

software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program"; below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any

warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you

cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

OpenSSL License

Copyright (c) 1998-2007 The OpenSSL Project. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:
"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"
4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.

6. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Copyright Statement

COPYRIGHT

No part of this publication or documentation accompanying this Product may be reproduced in any form or by any means or used to make any derivative work by any means including but not limited to by translation, transformation or adaptation without permission from Netronome Systems, Inc., as stipulated by the United States Copyright Act of 1976. Contents are subject to change without prior notice.

Copyright © 2013 Netronome Systems, Inc. All rights reserved.

Limited Warranty

WARRANTY

Netronome warrants that any media on which this documentation is provided will be free from defects in materials and workmanship under normal use for a period of ninety (90) days from the date of shipment. If a defect in any such media should occur during this 90-day period, the media may be returned to Netronome for a replacement.

NETRONOME DOES NOT WARRANT THAT THE DOCUMENTATION SHALL BE ERROR-FREE. THIS LIMITED WARRANTY SHALL NOT APPLY IF THE DOCUMENTATION OR MEDIA HAS BEEN (I) ALTERED OR MODIFIED; (II) SUBJECTED TO NEGLIGENCE, COMPUTER OR ELECTRICAL MALFUNCTION;

OR (III) USED, ADJUSTED, OR INSTALLED OTHER THAN IN ACCORDANCE WITH INSTRUCTIONS FURNISHED BY NETRONOME OR IN AN ENVIRONMENT OTHER THAN THAT INTENDED OR RECOMMENDED BY NETRONOME. EXCEPT FOR WARRANTIES SPECIFICALLY STATED IN THIS SECTION, NETRONOME HEREBY DISCLAIMS ALL EXPRESS OR IMPLIED WARRANTIES OF ANY KIND, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow the exclusion of implied warranties, so the above exclusion may not apply to some users of this documentation. This limited warranty gives users of this documentation, specific legal rights, and users of this documentation may also have other rights which vary from jurisdiction to jurisdiction.

Liability

Regardless of the form of any claim or action, Netronome's total liability to any user of this documentation for all occurrences combined, for claims, costs, damages or liability based on any cause whatsoever and arising from or in connection with this documentation shall not exceed the purchase price (without interest) paid by such user.

IN NO EVENT SHALL NETRONOME OR ANYONE ELSE WHO HAS BEEN INVOLVED IN THE CREATION, PRODUCTION, OR DELIVERY OF THE SSL INSPECTOR APPLIANCE BE LIABLE FOR ANY LOSS OF DATA, LOSS OF PROFITS OR LOSS OF USE OF THE SSL INSPECTOR APPLIANCE OR FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, EXEMPLARY, PUNITIVE, MULTIPLE OR OTHER DAMAGES, ARISING FROM OR IN CONNECTION WITH THE SSL INSPECTOR APPLIANCE EVEN IF NETRONOME HAS BEEN MADE AWARE OF THE POSSIBILITY OF SUCH DAMAGES.

IN NO EVENT SHALL NETRONOME OR ANYONE ELSE WHO HAS BEEN INVOLVED IN THE CREATION, PRODUCTION, OR DELIVERY OF THE SSL INSPECTOR APPLIANCE BE LIABLE TO ANYONE FOR ANY CLAIMS, COSTS, DAMAGES OR LIABILITIES CAUSED BY IMPROPER USE OF THE SSL INSPECTOR APPLIANCE OR USE WHERE ANY PARTY HAS SUBSTITUTED PROCEDURES NOT SPECIFIED BY NETRONOME.

Disclaimer of Damages

SOME STATES AND COUNTRIES, INCLUDING MEMBER COUNTRIES OF THE EUROPEAN ECONOMIC AREA, DO NOT ALLOW THE LIMITATION OR EXCLUSION OF LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES SO THE BELOW LIMITATION OR EXCLUSION MAY NOT APPLY TO YOU.

TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW AND REGARDLESS OF WHETHER ANY REMEDY SET FORTH HEREIN FAILS OF ITS

ESSENTIAL PURPOSE, IN NO EVENT WILL NETRONOME OR ITS LICENSORS OR PARTNERS, BE LIABLE TO YOU FOR ANY SPECIAL, CONSEQUENTIAL, INDIRECT OR SIMILAR DAMAGES, INCLUDING ANY LOST PROFITS OR LOST DATA ARISING OUT OF THE USE OR INABILITY TO USE THE APPLIANCE, EVEN IF NETRONOME HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

IN NO CASE SHALL NETRONOME'S OR ITS LICENSORS OR PARTNER'S LIABILITY EXCEED THE PURCHASE PRICE FOR THE APPLIANCE.

THE DISCLAIMERS AND LIMITATIONS SET FORTH ABOVE WILL APPLY REGARDLESS OF WHETHER YOU ACCEPT THE APPLIANCE OR ITS ASSOCIATED SOFTWARE.

Export Regulations

You agree to comply strictly with all applicable export control regulations and laws, including the US Export Administration Act and its associated regulations and acknowledge Your responsibility to obtain licenses as required to export, re-export or import the Appliance. Export or re-export of the Appliance to the following countries is prohibited: Cuba, North Korea, Iran, Iraq, Libya, Syria and Sudan.

General

If you are located in North America or Latin America, this Agreement will be governed by the laws of the state of California, United States of America. Otherwise, this Agreement will be governed by the laws of England. This Agreement and any related License module is the entire between you and Netronome Systems, Inc. relating to the Appliance and (i) supersedes all prior or contemporaneous oral or written communications, proposals and representations with respect to its subject matter, and (ii) prevails over any conflicting or additional terms of any quote, order, acknowledgement or similar communications between the parties. This agreement may only be modified by a License Module or by a written document which has been signed by both you and Netronome. This agreement shall terminate upon Your breach of any term contained herein and You shall cease use of the Appliance and its associated software and shall return the Appliance to Netronome. The disclaimers of warranties and damages and limitations on liabilities shall survive termination. Should you have any questions concerning this agreement, please write:

Netronome Customer Service
3159 Unionville Drive, Suite 100
Cranberry Township, PA 16066
USA

Excluded Software

The excluded software consists of the open source code software known as Linux included in the Appliance. All Excluded Software is licensed under the GNU General Public License, Version 2, June 1991, a copy of which is included in this document. The license entitles You to receive a copy of the source code for the Linux only upon request at a nominal charge. If you are interested in obtaining a copy of such source code, please contact Netronome Systems, Inc. for further information.

Chapter 9

Technical Support

To obtain additional information or to provide feedback, please email support@sourcefire.com or contact the nearest Sourcefire technical support representative.

Visit <https://support.sourcefire.com> to download the latest documentation and software, access the knowledge base, or log a support ticket.