Cisco SSL Appliance Administration & Deployment Guide

Version 3.9
November 20, 2015

Cisco Systems, Inc.
www.cisco.com

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

Introduction

The following conventions are used throughout this document.

Note This style indicates a "note" providing additional information that the reader might be interested in.

Caution This style indicates a "warning" providing additional information that the reader needs to pay attention to.

Elements you see on the WebUI such as the names of screens, fields, and options **appear in this typeface**.

This guide covers all SSL Appliances and their variations. Where differences exist, the models are called out.

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 SSL Appliance is capable of inspecting both types of traffic.

The SSL Appliance software is subject to licensing by Cisco. See System Status, page 4-6 of this document for details on licensing.

Note The act of “inspecting” SSL traffic might be subject to corporate policy guidelines and/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.

SSL Inspection Overview

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. SSL 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 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 e-mail).

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

SSL Inspection enables existing security and network appliances to access the plaintext within SSL flows thereby enabling the security appliance to do its job, even with SSL encrypted traffic. Unmodified applications running on devices attached to the SSL Appliance gain visibility into the content of the SSL traffic. 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 in order 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 resign 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 "In-line" or can be connected to a network span or tap port.

Note SSL Inspection using "certificate resign" and SSL policy enforcement can only be done if the SSL Inspecting device is connected "inline" in the network.

Note Only "known server key" mode can be used to inspect SSL traffic when the inspecting device is connected to a network tap. Inspection is not possible if the session uses Diffie-Hellman or Elliptic Curve Diffie-Hellman for key exchange.

SSL inspection 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

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

Note The SSL Appliance and the associated security appliance(s) that it is enabled to “inspect” traffic should all be located in a physically secure environment in order to prevent unauthorized access to the decrypted SSL traffic.

Product Overview

The Cisco SSL Appliance is a high performance transparent proxy for Secure Socket Layer (SSL) network communications. It enables a variety of applications to access the plaintext (that is, the 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 lets network appliances be deployed with highly granular flow analysis while maintaining line rate performance.

Cisco's SSL Appliance products provide two main functions:

- Enabling other security appliances to see a non encrypted version of SSL traffic that is crossing the network. This is called SSL Inspection, as the security appliance is able to inspect the decrypted traffic for possible threats: something it cannot do when it sees encrypted traffic.
- Acting 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 TCP destination port number being used by a session to determine if it is using SSL or not. The SSL Appliance uses deep packet inspection (DPI) to identify SSL flows. This ensures that it can find and inspect 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 multi-Gigabit/s 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 SSL inspection. 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 resign 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.

There are three basic connectivity modes that define how the SSL 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 “In-line” or connected to a network span or tap port.

Note SSL Inspection using “certificate resign” and SSL policy enforcement can only be done if the SSL Appliance is connected “in-line” in the network.

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

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.

Note The SSL Appliance and the associated security appliance(s) that it enabled to “inspect” traffic should all be located in a physically secure environment in order to prevent unauthorized access to the decrypted SSL traffic.

Note The act of “inspecting” SSL traffic might be subject to corporate policy guidelines and/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 provides 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

All non SSL traffic flows are “cut through” (forwarded directly from port to port) by the embedded flow processor(s), minimizing latency for traffic such as 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 / 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, such as web (HTTPS), e-mail 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 when 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 might 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 in-line) 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.

Security Best Practices

This section presents recommendations from Cisco on best practices for running the SSL Appliance securely.

Master Key Storage

The SSL Appliance stores security sensitive configuration in secure storage, which is encrypted with keys derived from a master key. The appliance obtains the master key when it boots up in order to unlock the secure storage and access appliance configuration. See the Getting Started Guide for your SSL Appliance for information about creating and storing the master key.

The master key storage location can be configured during the initial configuration (bootstrap) of the appliance:

- Internal: the master key is stored on an internal persistent storage device
- USB: the master key is stored externally on a removable USB persistent storage device

In situations where having the ability to power on the device should not place the device in an operational state, Cisco recommends that the appliance be configured to store the master key on a removable USB storage device. With this option, the SSL Appliance will require the USB storage device to be physically plugged in before it can boot and process network traffic. When not plugged into the appliance, the USB storage device should be stored in a secure location.

Users can also configure the SSL Appliance to protect the master key with a master key password. Cisco recommends that passwords of 8 or more characters be used to protect from brute-force attacks. To configure a master key password, select the “Keypad Password” master key protection option during the appliance bootstrap phase. This option will require a user to physically enter the master key password through the front panel keypad every time the appliance is powered up or restarted.

Creating User Accounts

The SSL Appliance uses role-based authentication with separation of duties, which prevent authenticated users from performing undesired actions. For example:

- The Manage Appliance role is for system and network administrators, who need to manage the system and network settings.
- The Manage Policy role is for policy and compliance administrators, who need to manage rules for processing and decrypting SSL traffic.
- The Manage PKI role is for PKI administrators who need to manage the cryptographic keys, certificates and CRLs on the appliance.

- The Auditor role is for auditors, who need to view the appliance configuration and log files, but cannot modify any settings.

Cisco recommends creating different user accounts with different user roles to implement a level of privilege separation appropriate for the security policies and requirements set by the organization deploying the SSL Appliance. Elevated privilege separation would prevent authenticated users from performing unauthorized functions, e.g. an auditor from modifying the decryption policy. For more information about the functions authorized for each user role, refer to Users, page 6-65 in this document.

Management Interfaces Access Control

Cisco recommends that the SSL Appliance be deployed in a secured network that appropriately restricts access to the appliance management network interfaces. This will prevent non-authorized users from accessing the appliance management interfaces – WebUI, command line diagnostics (CLD) interface, and SNMP interface. Network administrators can use firewall rules to restrict access to the SSL Appliance's IP address. Authenticated users with the Manage Appliance role can also use the Access Control Lists feature to restrict incoming WebUI, CLD, and SNMP connections.

The SSL Appliance intercepts and decrypts network traffic on dedicated physical network interfaces, which are physically separate from the appliance management interfaces. The data interfaces do not provide management access to the SSL Appliance. Therefore, the data interfaces need not be subject to the same access control restrictions as the management network interfaces.

Secure SNMP Monitoring

The SSL Appliance supports the more secure SNMP version 3, which supports authentication and encryption for SNMP monitoring. Even though SNMP versions 1 and 2c are also supported, Cisco recommends keeping them disabled to prevent non-authorized users from monitoring the appliance. Cisco also recommends the default options of using AES for encryption and SHA for authentication for SNMP version 3.

Certificate Resigning Using Secure Key Sizes

When the SSL Appliance policy decrypts an SSL flow using a "Decrypt (Resign Certificate)" rule, the appliance will modify and resign the flow's server certificate. It signs the certificate using a resigning certificate authority configured through the WebUI. Resigning a server certificate with an insecure certificate authority key is a security vulnerability because it allows a downstream attacker to perform a man-in-the-middle attack and decrypt the encrypted TLS payload. When configuring resigning certificate authorities in the PKI store, Cisco recommends that customers use RSA keys of size 2048 bits or higher, or Elliptic Curve keys on curves of size 224 bits or higher.

HSM Authentication Using Secure Key Sizes

When using an HSM appliance for certificate resigning, the SSL Appliance authenticates to the HSM using a TLS client certificate and key. Using insecure key sizes for client authentication may allow an attacker to impersonate an SSLV appliance and get a malicious certificate signed by the HSM. Clients are typically configured to trust the HSM signing key certificate, allowing the attacker to stage a man-in-the-middle attack. Cisco recommends that customers configure client authentication to HSM appliances to use RSA client certificates and keys of size 2048 bits or higher.

SSL Appliance WebUI Keys and Certificates Using Secure Key Sizes

The WebUI is only accessible over encrypted and authenticated TLS connections. The WebUI server is configured with a private key and certificate for TLS authentication. Using insecure key sizes may allow an attacker to impersonate the WebUI server and cause appliance administrators to inadvertently reveal sensitive password information. Cisco recommends that customers use RSA keys of size 2048 bits or higher for the appliance WebUI.

Wiping Appliance Hard Disk Before RMA

The SSL Appliance stores configuration data, log files and statistics on internal persistent storage. Before returning an already configured or deployed appliance to Cisco, Cisco recommends that customers erase data on internal persistent storage devices. Erasing persistent data is done using the factory reinstall operation:

Step 1 Connect a serial console, or monitor and keyboard to the appliance.

Step 2 Power up or restart the appliance.

Step 3 Choose the Factory re-install option from the bootloader menu that appears at startup.

Step 4 Wait for the appliance to clear the contents of the persistent storage devices, perform the factory reinstall operation, and restart itself.

After the appliance restarts, user data will be zeroized (overwritten with zeros).

Note The zeroization procedure used by the “Factory re-install” option may not purge all sensitive data from the internal storage devices.

If the SSL Appliance was previously configured to store the master key on a USB persistent storage device, the storage device is now unusable and all data on it should be erased.

System Behavior & Deployment Examples

This section describes the functions performed by the SSL Appliance, its behavior, and its interaction with attached devices. For details on how to setup and configure the SSL Appliance, refer to Initial Configuration and Setup, page 4-1 and User Interface Overview, page 6-1. The *Getting Started Guide* provides detailed initial setup and connection information.

Transparent SSL Decryption / Encryption

The main function of the 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 device(s) for processing or analysis. As the plaintext data stream is repackaged as a valid TCP stream, applications that are hosted on the attached device(s) do not need to be modified to process the received plaintext stream.

- The SSL Appliance provides SSL Inspection capabilities to existing devices.

The collection of SSL Appliance interfaces that are used to connect to the network carrying the traffic that is being inspected and to the attached appliances that are processing the traffic is called a “segment”. Depending on how the appliance is connected, and on how many attached appliances are connected, a segment may contain up to 8 interfaces.

When used in Active-Inline (AI) mode or Passive-Inline (PI) 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. Unlike a non transparent proxy which requires that client machines are configured to send traffic to the IP address associated with the proxy there are no changes required to clients or other network equipment 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 it can be deployed without renumbering the existing IP network. In most cases no network topology changes whatsoever are required.
- If used in Passive-Tap (PT) 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 appliance(s).

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, such as HTTP, SMTP, POP3, IMAP, and many other proprietary protocols. 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 e-mail protocols (SMTP, POP3 and IMAP).

- The SSL Appliance can decrypt most SSL 3.0, TLS 1.0, TLS 1.1 and TLS 1.2 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 does 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 for inspecting 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 resign 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 these methods can be used when the SSL Appliance is operating in Active-Inline (Active-Inline Mode, page 3-11) or Passive-Inline (Passive-Inline Mode, page 3-10) mode but only the "known server key" method can be used if the SSL Appliance is operating in Passive-Tap Mode, page 3-8.

Note The method 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.

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 Policy Rulesets, page 3-15.

Known Server Key Method

The next figure 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 that you want to inspect traffic to. The SSL Appliance can use the key/certificate from a specific server to decrypt SSL sessions established with that server. If the private key only mode is being used, then references to key and certificate in the rest of this section should be taken to mean only the private key.

This method can only be used where the SSL Appliance administrator has access to the server private key and certificate information; this is normally only the case 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.

Known Server Key Passive-Tap

The simplest example of known server key mode is illustrated. You can see that the client is sending "abc" to the server, and this 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, it decrypts this to get the plaintext "abc."

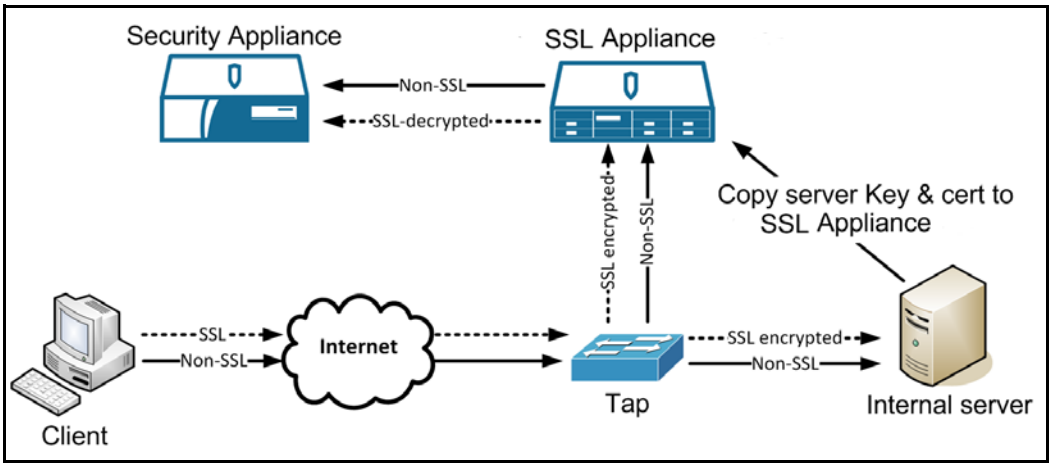

In this example, the SSL Appliance is not a "Man In The Middle" (MITM) of the SSL session. It is simply receiving a copy of the encrypted data, and decrypting it using the server private key and certificate that it has copies of.

The fact that in passive-tap mode the SSL Appliance is not a MITM for the SSL session is important, as 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. In order 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.

Known Server Key Passive-Inline

In Passive-Inline mode, the SSL Appliance is a MITM as the traffic between client and server passes through the SSL Appliance.

An important point to note here is that 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 can decrypt 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".

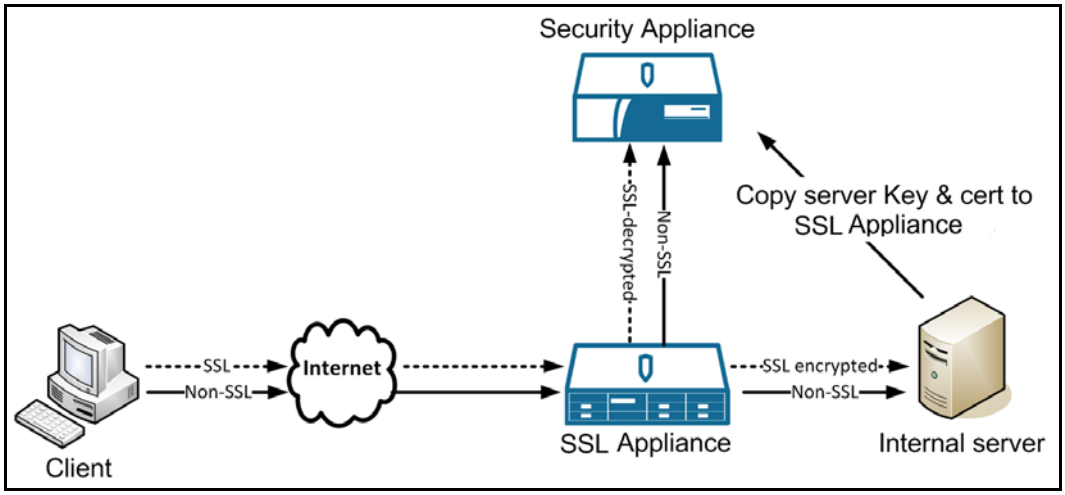

The encrypted traffic between the client and the SSL Appliance and between the SSL Appliance and the server is different, because the two SSL sessions have different cryptographic session details. If the session uses Diffie-Hellman for key exchange, the session details will be different for the two SSL sessions. 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.

Traffic to many different SSL servers with different SSL server certificates can be inspected by a single SSL Appliance.

Certificate Resigning Method

Certificate resign is used when it is impossible to obtain a copy of the SSL server's private key and certificate, which 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 will need to be inspected using certificate resign.

Note In order to use certificate resign, the SSL Appliance must be a MITM which means this mechanism cannot be used if the SSL Appliance is connected in Passive-Tap mode.

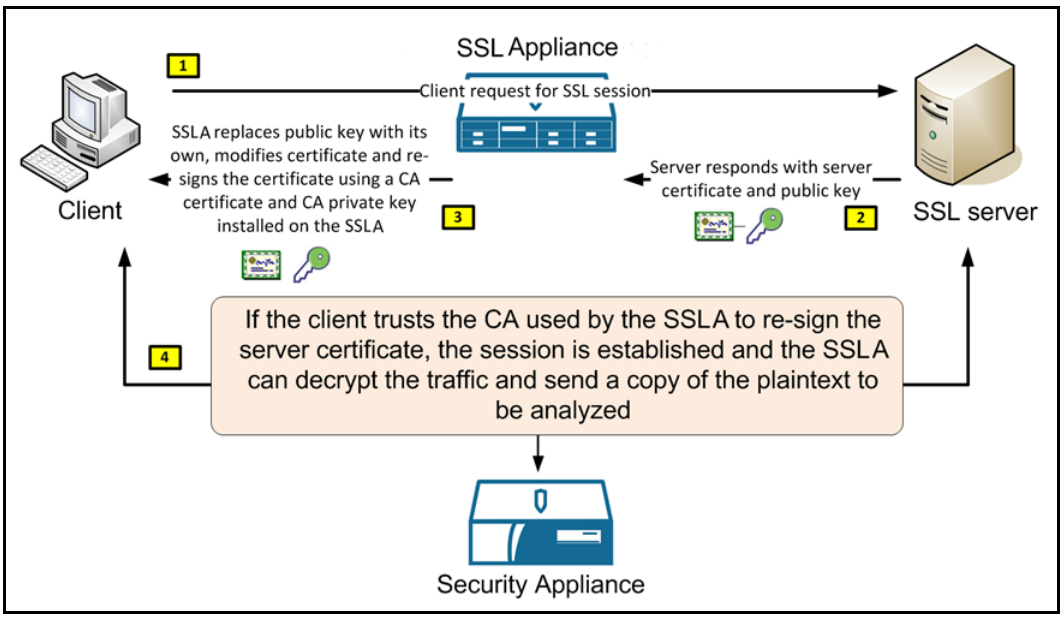

The client initiates an SSL session to the server and the server responds by sending it's SSL server certificate to the client. As all traffic between client and server passes through the SSL Appliance it can detect and intercept the server certificate.

Once the SSL Appliance has intercepted the server certificate, it replaces the servers public keys with its own public keys and modifies the Certificate Revocation List (CRL) details in the server certificate. Having modified the server certificate, the SSL Appliance then resigns the server certificate using a Certificate Authority (CA) certificate and CA private key that is installed in the SSL Appliance.

The resigned server certificate is then sent over the network to the client. If the client trusts the CA that was used to sign the server certificate it receives it will not generate any warnings. As 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 resign is used the two SSL sessions will always have different cryptographic session details and the SSL Appliance will have to re-encrypt the plaintext before sending it back to the network.

As noted, the client must trust the CA used to resign the server certificate; otherwise it will generate warnings indicating that the SSL session should not be trusted. In order to ensure that the client does trust the CA used by the SSL Appliance, there are two approaches that can be taken.

- The SSL Appliance can generate a CA certificate and keys internally and use these to resign server certificates. The CA certificate which includes the CA public key can be exported from the SSL Appliance, and then imported into the trusted CA store on the client; you only have to do this once.
- If the SSL Appliance is deployed in a network that already has a private public key infrastructure (PKI), this can be used to issue an intermediate CA certificate and keys which can be loaded into the SSL Appliance. As 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.

Use of EC Resigning CAs

Certificate authorities may sign server certificates with either RSA or Elliptical Curve keys.

If the system tries to use certificate resign to inspect a flow that has a server certificate signed by a CA using EC keys, and it resigns with an resigning CA that uses RSA keys, it won't work. The CA used to resign the server certificate must use the same type of key as the original CA.

Hence, the SSL Appliance must have two internal resigning CAs on the appliance, one that uses RSA keys, and another using EC keys. You can create or load keys that use either RSA or EC keys for use in resigning server certificates.

In the SSL inspection rules, you can specify an external resigning CA that uses RSA keys, and another that uses EC keys. If a CA using EC keys is not present, a flow with an EC signed server certificate will not match the rule, and will normally be cut through.

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, so some organizations might 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 Policy Rulesets, page 3-15).

If SSL connections to servers using self-signed certificates are allowed, the SSL Appliance can inspect the traffic two ways.

- Resign the certificate the same way a none self-signed certificate is resigned; see Certificate Resigning Method, page 3-4. This method is used if **Decrypt (Resign)** mode is chosen.
- The second method involves the self-signed certificate information (that is, 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 used.

If 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 might be deployed in networks that already have an SSL proxy device in place that is inspecting some of the outgoing SSL traffic using certificate resign. The SSL Appliance would typically be deployed in order to allow other security appliances to view inspected traffic in addition to the existing proxy device that might 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.

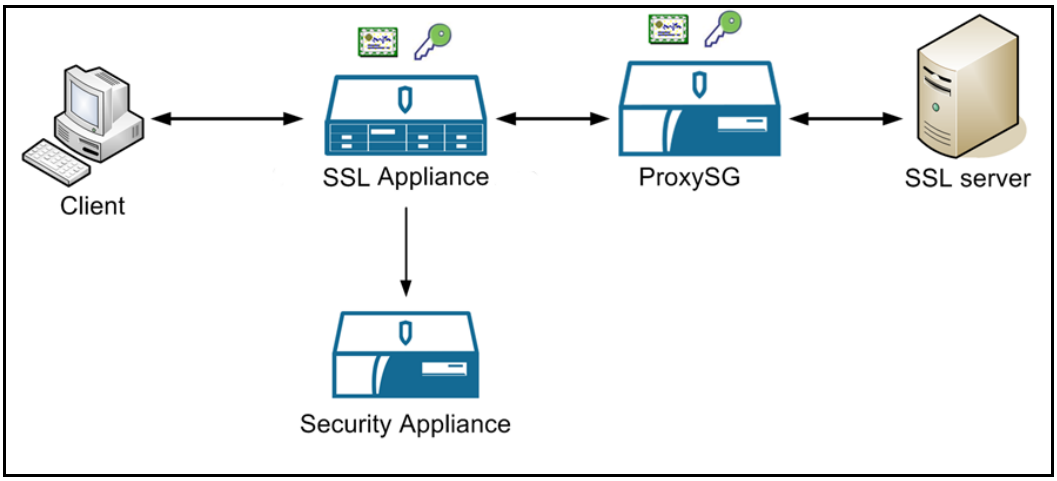

In a Passive-Inline mode using certificate resign cooperative configuration, both the existing SSL proxy and the SSL Appliance are MITM devices. The existing proxy resigns the original server certificate and then the SSL Appliance resigns the modified server certificate it receives. If the SSL Appliance doesn't trust the proxy's resigning certificate, it will mark the flow with the "invalid issuer" status. It will however resign the flow. The flow will not be resigned only if the default rule configuration is modified such that the rule doesn't apply to flows with "invalid issuer" certificate status.

No matter how many devices are resigning the server certificate, the client only sees the last CA which signed the certificate. This last signing certificate needs to be added to the client store.

Mark SSL Plaintext

The generated flow containing plaintext obtained from inspected SSL traffic can optionally 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, as the SSL Appliance needs to 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 if VLAN marking is being used.

Deployment Modes

This section provides details on how the SSL Appliance can be deployed in a network and how it operates in each of the deployment modes. The deployment mode is configured for a segment, each segment uses a number of network interfaces on the SSL Appliance. There might be multiple 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.

Before looking at the deployment modes in more detail we need to define some terminology that 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 which is dealing with inspected traffic from the SSL Appliance.
- Copy port: A network interface connected to a secondary passive appliance that is receiving a copy of the inspected traffic.
- Aggregation port: A network interface providing a connection to an additional network tap, so that a segment can receive traffic from more than one network tap.
- Symmetric traffic: Traffic where packets for both directions of a network flow are seen on the same network interface on the SSL Appliance.
- Asymmetric traffic: Traffic where the packets for both directions of a network flow are seen on different network interfaces on the SSL Appliance.
- Active-active: An HA deployment scenario where packets on a given flow might be 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 might 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 Example Passive-Tap Mode Inspection, page 4-12, Example Passive-Inline Mode Inspection, page 4-20, Example Active-Inline Mode Inspection, page 4-25, and Configure Receive Interfaces with Segments, page 6-23.

Note 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 from the set of interfaces that are not currently in use by other, already active, segments.

Segment Elements

Segment configuration can be considered to have five elements; not all of these elements apply to a given segment:

- The network interfaces connecting traffic to the SSL Appliance. In a passive-tap mode, the minimum number of such interfaces is one. In an in-line mode, the minimum number is two, as the SSL Appliance is a bump-in-the-wire.
- Whether the traffic being inspected is symmetric or asymmetric. If the traffic is asymmetric, more network interfaces are 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.

- Whether there is an active appliance connected to the SSL Appliance. An active appliance requires a minimum of two interfaces connecting it to the SSL Appliance.
- Whether there are any passive appliances connected to the SSL Appliance. A passive appliance requires a minimum of one interface connecting it to the SSL Appliance.
- Whether there is more than one passive appliance connected to the SSL Appliance. If more than one passive appliance is connected, then decide if all traffic should be copied to each passive appliance, or if it should be load balanced between the passive appliances.

Passive-Tap Mode

This section provides details on the different Passive-Tap modes of operation supported by the SSL Appliance. Passive-Tap mode connectivity options fall into three groups based on:

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

Note Only known server key decryption can be used when the SSL Appliance is deployed in Passive-Tap mode.

Note If Diffie-Hellman is used for key exchange then the SSL Appliance will be unable to decrypt the flow using the know server key methods when it is connected in Passive-Tap mode.

One common use for Passive-Tap mode is to connect an SSL Appliance to the network configured not to inspect any SSL traffic but with the SSL Session Log enabled. This is a quick way to collect 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 SSL Session Log provides a detailed picture of the SSL traffic in the network and can be used to plan what traffic needs to be inspected and how the SSL Appliance needs to be connected to the network to achieve this.

Passive-Tap Symmetric

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

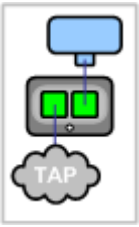

The simplest passive-tap deployment (passive-tap symmetric) has the SSL Appliance connected to a tap that delivers symmetric traffic to the SSL Appliance over a single network interface. The inspected traffic is then sent to a single passive appliance as symmetric traffic over a single network Interface.

There are two common deployments that 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).

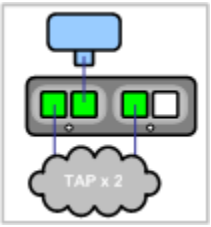

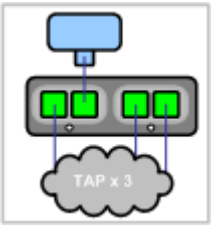

Note If two tap ports are being used in aggregation mode and are connected to interfaces that share fail-to-wire hardware then whenever the FTW is active the two taps will be connected to each other. You are advised to ensure that this will not cause problems for the tap ports or the network.

Any of the above symmetric Passive-Tap modes can be configured to use an additional two interfaces (copy ports) for connection to additional attached passive appliances.

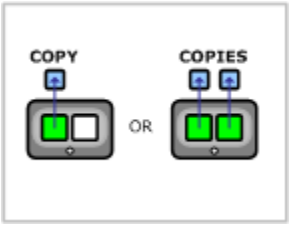

If a single copy port is used, it feeds a copy of the symmetric traffic from the SSL Appliance to the first passive appliance. If two copy ports are used, these can be used to either:

- 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

Passive-Tap Asymmetric

Passive-tap mode also supports inspection of asymmetric traffic.

Copy options are available for this asymmetric Passive-Tap mode of operation.

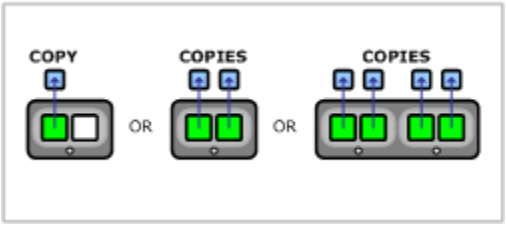

If no copy ports are used then a single passive appliance will receive the asymmetric traffic from the SSL Appliance over the two device ports.

If a single copy port is used then it feeds a symmetric copy of the asymmetric traffic from the SSL Appliance to a second passive appliance. If two interfaces are used these can be used to either:

- 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 these can be used to either:

- 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

Passive-Inline Mode

This section provides details on the different Passive-Inline modes of operation supported by the SSL Appliance. Passive-Inline mode connectivity options fall into two groups based on:

- Is the SSL Appliance connected inline on a network segment that carries traffic for both directions of a flow? This is a symmetric traffic case.
- Is the SSL Appliance connected inline on two network segments with packets for a given flow potentially being present on one or other segment? This is an asymmetric traffic case.

Note If the SSL Appliance is being deployed in a network using an active-active HA architecture, treat this as an asymmetric traffic case. The SSL Appliance can be configured as an in-line 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.

Symmetric Passive-Inline

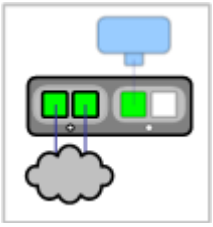

Copy port options are available for symmetric Passive-Inline mode. In Passive-Inline mode there are no device ports configured as part of the initial segment configuration, so all attached appliances are connected to copy ports.

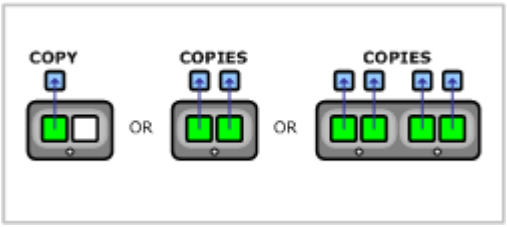

If a single copy port interface is used, it will feed a symmetric copy of the symmetric traffic from the SSL Appliance to the first passive appliance. If two interfaces are used, they can either

- 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 interfaces are used, they can be used to either:

- 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

Asymmetric Passive-Inline

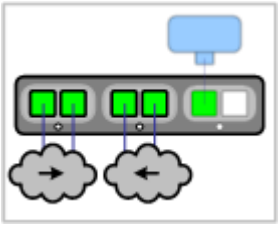

Use Passive-Inline mode to inspect asymmetric traffic.

Copy port options are available.

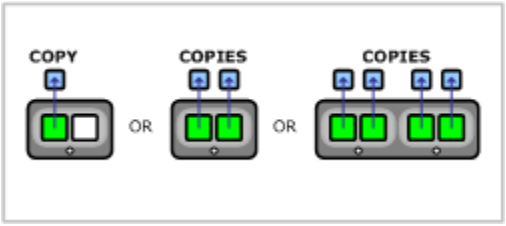

In Passive-Inline mode there are no device ports configured as part of the initial segment configuration so all attached appliances are connected to copy ports.

If a single copy port interface is used, it will feed a symmetric copy of the symmetric traffic from the SSL Appliance to the first passive appliance. If two interfaces are used, they can either

- 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 interfaces are used, they can be used to either:

- 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 options fall into two groups based on:

- Is the SSL Appliance connected inline on a network segment that carries traffic for both directions of a flow? This is a symmetric traffic case.
- Is the SSL Appliance connected inline on two network segments with packets for a given flow potentially being present on one or other segment? This is an asymmetric traffic case.

Note If the SSL Appliance is being deployed in a network using an active-active HA architecture this can be treated as an asymmetric traffic case. The SSL Appliance can be configured as an in-line 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, they way in which the active appliance is connected determines how traffic flows in the event of a failure of the SSL Appliance.

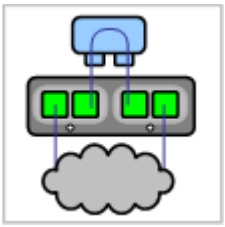

Fail To Appliance (FTA) mode results in traffic flowing through the attached active appliance in the event of failure.

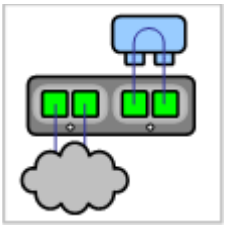

Fail To Network (FTN) mode results in traffic bypassing the active appliance in the event of failure.

Copy port options are available in Active-Inline mode.

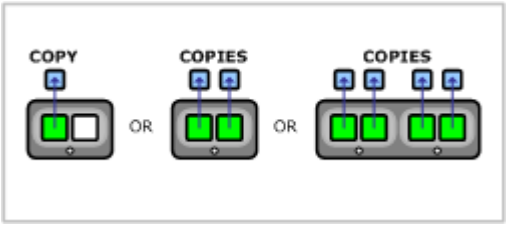

If a single copy port interface is used, it will feed a symmetric copy of the symmetric traffic from the SSL Appliance to the first passive appliance. If two interfaces are used, they can either

- 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 interfaces are used, they can be used to either:

- 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

Note An active security appliance which is attached to an Active-Inline segment must be transparent; it may not terminate or re-originated TCP, or modify packet headers. For example, it may not run a process such as Network Address Translation (NAT) which changes header details. Intrusion prevention systems (IPS), however, interact correctly with the SSL Appliance, as they do not change headers.

Asymmetric Active-Inline

Active-Inline asymmetric may FTA, or FTN.

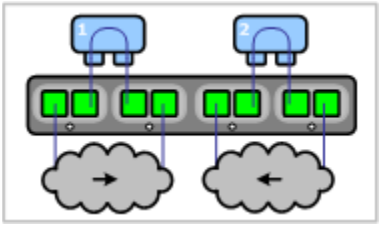

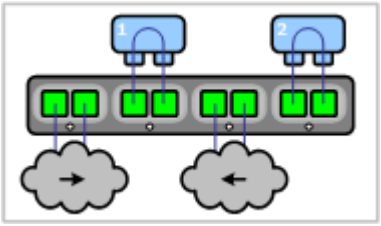

Copy port options are the same as for symmetrical mode.

Policies

Policies in the SSL Appliance are composed of three elements:

- Lists; referenced by rules in rulesets
- Segments; grouping of networking interfaces receiving traffic, set the ruleset to use and the deployment mode
- Rulesets; containing the rules which control how SSL traffic is handled; associated with one or more segments in use.

Lists are used to collect multiple items of the same type of information so that a single ruleset can point to the list and is applied whenever any of the items in the list are true. For example, a list might contain 20 different Subject/Domain Names (S/DN) that occur in the server certificates from 20 different sites, a policy that is configured to "inspect" traffic when it detects a particular Subject/Domain Name can point to the list instead of just indicating a single Domain Name in the policy. This allows a single policy entry to apply to all 20 different sites and means that additional sites can be added (by editing the list) without needing to edit the ruleset.

A segment is a grouping of interfaces that receives a network feed; it tells the SSL Appliance which Ruleset to use and in what deployment mode to operate with that network feed, and how to distributed the decrypted SSL and other received traffic. 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 A could be using ruleset1 and segment 2 ruleset4 or both segments A and B could be using ruleset3. Inactive segments are not associated with physically 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, page 3-7. The system allocates 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 following table.

Table 3-1 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 Mode	Disabled	Save: • Disabled • Local Session Log Only • Local Session Log, All sessions to Remove Syslog • Local Session Log, Errors to Remote Syslog
Compression	Cut through	The 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
SSL2	Cut through	
Diffie-Hellman in Passive-Tap mode	Cut through	
Client Certificate	Reject	
Cipher suite (including Export)	Cut through	
Uncached	Cut through	
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 valid yet		
Self-signed		
Revoked		
Status Override Order	Rule over Segment	

Note The Session Log Mode must match the "Remote Logging" Options configuration in order to send the specified logs to a remote syslog server.

Policy Rulesets

A ruleset has a fixed set of options and a variable number of rules. A rule is used to match against a specific SSL flow or set of flows. The SSL Appliance can be very specific in matching a flow using a rule, be more general by using a list of rules, or be “generic” in matching all flows. Modify the parameters of a rule, and the structure of a ruleset to achieve the granularity you want. In the following tables any entry where the Default Setting field is empty means that the default setting is the “nothing is set” option.

The SSL Appliance extracts CN, Subject Alternative Name (SAN), and Server Name Indication (SNI) information from intercepted flows in order to deduce the SSL server domain name. The SSL flows are matched against rules using this process:

Step 1 The SSL Appliance policy rules support the following subject distinguished name (DN) attributes:

- CN: Common Name
- O: Organization
- OU: Organizational Unit
- C: Country

Step 2 **Subject/Domain Name** and **Subject/Domain Name List** match field entries without a prefix, as well as all **Domain Names List** match field entries, are treated as domain names, and are matched against the domain name deduced from the SSL flow. The rules match fields can contain "*" wild card characters, which will be expanded when matching. For example, a rule match field domain name "*.company.com" will match SSL flows with domain names.

- The SSL Appliance matches the SNI hostname from the SSL flow to the server certificate's subject CN and SAN entries. If a match is found, the SNI hostname is treated as the flow's domain name. If there is no SNI hostname in the flow, or if it does not match any subject CN or SAN entries, the union of all {subject CNS, SAN entries} is considered as possible domain names.
- The SSL Appliance matches the deduced domain name(s) to the domain name match fields in the rule match fields. If a domain name matches, the match field is considered to match.

The following table shows the basic set of policy options contained in a ruleset. A single ruleset can have one or more rules. The details relating to rules themselves are shown in more detail later in this section.

Table 3-2 Policy Ruleset Options

Item	Default Setting	Notes
Name		Identifies this ruleset
Default RSA Resigning Certificate Authority		Default RSA CA used for certificate resign
Default EC Resigning Certificate Authority		Default EC CA used for certificate resign
External Certificate Authorities	All external CAs	Can point to a custom list instead
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
Host Categorization IP Exclude List		IP list used to prevent Host Categorization lookup.
HSM Failure Action	Cut through	Action where the HSM resign operation has failed; also Reject, Drop

There are several different 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, and these fields are compared with the corresponding values in an SSL session to determine if the rule should be applied to the session or not.

Any match fields that are left empty are treated as matching any value for that field. The different rule types allow for a total of eight possible actions that can be taken if a rule is matched, these are listed in the following table.

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

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 to point to a specific item, and another field to point 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.

Note The **Subject/Domain Name, Subject/Domain Name List**, and **Domain Name List** are mutually exclusive.

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

The following table shows details for a Decrypt (Certificate and Key known) rule that triggers decryption using a known server key and certificate if the details in the server certificate for a session match the rule.

Table 3-3 Decrypt with Known Certificate and Key Rule Formats

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 mask so can specify subnet
Source IP List	↑	Name of list of source address/masks that is checked for a match
Destination IP	-	IP address and mask so can specify subnet
Destination IP List	↑	Name of list of destination address/masks that is checked for a match
Destination Port		Destination TCP port number
Host Categorization List		Name of Host Categorization List checked for a match.
Traffic Class Unconfigured		Policy not based on Traffic Classes

Table 3-3 Decrypt with Known Certificate and Key Rule Formats

Item	Default Setting	Notes
Traffic Class (Value|Mask)		Rules refer to these values only, if used
Traffic Class List		Name of Traffic Class list that is checked for a match

The following table shows details for a Replace Certificate and Key rule that triggers decryption using a certificate and key replacement method if the details in the server certificate for a session match the rule. 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.

Table 3-4 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
RSA Known Certificate with Key (to replace with)		Pointer to an RSA certificate and key that will be used to replace the certificate and key in the server certificate
EC Known Certificate with Key (to replace with)		Pointer to an EC 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/Domain Name	-	Subject/Domain names checked for a match; server domain names captured via CN, SAN, SNI fields.
Subject/Domain Name List	↑	List of Subject/Domain names checked for a match; server domain names captured via CN, SAN, SNI fields.
Domain Name List		List of Domain names checked for a match.
Issuer DN	-	Issuer Subject/Domain Names checked for a match.
Issuer DN List	↑	List of Issuer Subject/Domain Names checked for a match.
Source IP	-	IP address and mask so can specify subnet
Source IP List	↑	Name of list of source address/masks that is checked for a match
Destination IP	-	IP address and mask so can specify subnet
Destination IP List	↑	Name of list of destination address/masks that is checked for a match
Destination Port		Destination TCP port number
Certificate Status		Status of X.509 server certificate
Host Categorization List		Name of Host Categorization List checked for a match.
Traffic Class Unconfigured		Policy not based on Traffic Classes
Traffic Class (Value|Mask)		Rules refer to these values only, if used
Traffic Class List		Name of Traffic Class list that is checked for a match

The following table shows details for a Decrypt (Resign Certificate) rule that will trigger decryption using certificate resign if the details in the server certificate for a session match the rule. 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.

Table 3-5 Decrypt using Certificate Resign Format

Item	Default Setting	Notes
Decrypt (Resign Certificate)		Decrypt using certificate resign
Comment		Optional descriptive text
RSA resigning CA		Pointer to the resigning RSA CA that is used to resign the server certificate
EC resigning CA		Pointer to the resigning EC CA that is used to resign the server certificate
Cipher Suite list		List of cipher suites: can't 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/Domain Name	-	Subject/Domain names checked for a match; Server domain names captured via CN, SAN, SNI fields.
Subject/Domain Names List	↑	List of server Subject/Domain names checked for a match.
Domain Name List		List of Domain names checked for a match.
Issuer DN	-	Issuer Subject/Domain Names checked for a match
Issuer DN List	↑	Issuer Subject/Domain Names checked for a match
Source IP	?	IP address and mask so can specify subnet
Source IP List	↑	Name of list of source address/masks that is checked for a match
Destination IP	-	IP address and mask so can specify subnet
Destination IP List	↑	Name of list of destination address/masks that is checked for a match
Destination Port		Destination TCP port number
Certificate Status		Status of X.509 server certificate
Host Categorization List		Name of Host Categorization List checked for a match.
Traffic Class Unconfigured		Policy not based on Traffic Classes
Traffic Class (Value|Mask)		Rules refer to these values only, if used
Traffic Class List		Name of Traffic Class list that is checked for a match

The following table shows details for a Decrypt (Anonymous Diffie-Hellman) rule that will trigger decryption if the details in the server certificate for a session match the rule. 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.

Table 3-6 Decrypt Anonymous Diffie_He;llman Format

Item	Default Setting	Notes
Decrypt (Anonymous Diffie-Hellman)		Decrypt Anonymous Diffie-Hellman session
Comment		Optional descriptive text
Source IP	-	IP address and mask so can specify subnet
Source IP List	↑	Name of list of source address/masks that is checked for a match
Destination IP	-	IP address and mask so can specify subnet
Destination IP List	↑	Name of list of destination address/masks that is checked for a match
Destination Port		Destination TCP port number
Host Categorization List		Name of Host Categorization List checked for a match.
Traffic Class Unconfigured		Policy not based on Traffic Classes
Traffic Class (Value|Mask)		Rules refer to these values only, if used
Traffic Class List		Name of Traffic Class list that is checked for a match

The following 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. 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.

Table 3-7 Non-Decryption Format Rules

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: can include Anonymous Diffie-Hellman cipher suites
Trusted Certificate	-	Certificate that is checked for a match
Trusted Certificates	↑	List of Certificates that are checked for a match
Subject/Domain Name	-	Subject/Domain names checked for a match; server domain names captured via CN, SAN, SNI fields.
Subject/Domain Name List	↑	List of Subject/Domain names checked for a match; server domain names captured via CN, SAN, SNI fields.
Domain Name List	↑	List of Domain names checked for a match.
Issuer Domain Name	-	Issuer Subject/Domain Names checked for a match.
Issuer Domain Name List	↑	List of Issuer Subject/Domain Names checked for a match.
Source IP	-	IP address and mask so can specify subnet
Source IP List	↑	Name of list of source address/masks that is checked for a match
Destination IP	-	IP address and mask so can specify subnet
Destination IP List	↑	Name of list of destination address/masks that is checked for a match
Destination Port		Destination IP port number

Table 3-7 Non-Decryption Format Rules

Item	Default Setting	Notes
Certificate Status		Status of X.509 server certificate
Host Categorization List		Name of Host Categorization List checked for a match.
Traffic Class Unconfigured		Policy not based on Traffic Classes
Traffic Class (Value|Mask)		Rules refer to these values only, if used
Traffic Class List		Name of Traffic Class list that is checked for a match

Lists

Lists can be referenced by rules in rulesets, and allow a single rule to be applied to more than one flow, as any flow that matches an entry in the list will trigger the rule action.

For each type of PKI list, the system creates a default list that is read only, and includes all items of that type present in the system. The default lists have names that begin with "all-" apart from the list of unsupported sites. User created custom lists are subsets of the default lists. The following table shows the default set of lists.

Name	Contains
all-external-certificate-authorities	All trusted external CAs
all-certificate-revocation-lists	All pointers to Certificate Revocation Lists
all-known-certificates	All known server certificates
all-known-certificates-with-keys	All known server private key/certificates
sslng-unsupported-sites	Sites it is not possible to inspect SSL sessions to

New keys or certificates are always imported 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, the system can contain lists of:

- Subject/Domain Names: Values without explicit distinguished name attribute types are considered domain names; the domain name values are matched against the SNI hostname, the subject Common Names (CNs), and the SAN DNS/IP entries. This includes the sslng-unsupported-sites list shown in the previous table.

Note Imported pre-3.7 policies using Distinguished Names lists are converted into Subject/Domain Names lists.

- Domain Names: Efficiently match SSL Appliance rules against website categories consisting of thousands of Domain Names.

Note Imported pre-3.7 policies using Common Names lists are converted into Domain Names lists.

- Cipher Suites
- IP addresses

The lists of Domain Names and lists of IP addresses are optimized to deal with large numbers of entries in the list as in some circumstances they might be configured with large numbers of entries.

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. Thus 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 Reject 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 transmits 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 round 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 then 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, so the SSL Appliance only generates 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 being 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.
- Decryption errors on a flow that is modified (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 a RST in both directions on the original SSL flow in order to kill the flow. If the active appliance generates a RST itself on the generated TCP flow then this will be detected by the SSL Appliance, and will trigger a 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): this is associated with a segment

- Software failure (data-plane): this is associated with the device

A segment is configured to operate in normal mode or High Availability (HA) mode. The failure actions taken by the device will differ depending on whether the segment is configured for HA mode or not. HA mode is not relevant if a segment is operating in Passive-Tap mode, so HA mode can only be configured 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, in order 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 it is operating in HA mode or not. Configuring HA mode enables the failure mode filter which is otherwise inactive.

When not operating in HA mode the failure of a link that is one of the links being used by the segment only 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 red in the status bar at the bottom of the screen
- The **Network** status indicator will change to red in the status bar at the bottom of the screen
- The event will be logged in the **System Log**
- If the link is part of the bump in the wire for an in-line segment or is the link to the network tap in PT mode, detection and inspection of SSL traffic will cease
- If the link is a link to an attached passive appliance, SSL detection and inspection will continue even though at least one of the attached passive appliances is no longer receiving the inspected traffic

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

- If the segment is Passive-Inline then failure of any segment interface will force all the network facing interfaces in the segment down
- If the segment is Active-Inline then 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 is configurable: either by manual reset from the WebUI or by auto recovery when the fault that triggered the failure is removed

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 and 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 then it triggers the failure mechanism.

The failure mode that becomes active when a failure occurs is configured per segment so a failure might 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 Configure Receive Interfaces with Segments, page 6-23 or 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

Modes that invoke Fail-to-wire cause the hardware mechanisms in the hardware FTM mechanisms 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 will be processed as link failures have priority over software failures.

Internally the system generates a recovery event once the issues that caused the software failure have been removed and all run-time tests have succeeded. Automatic recovery will occur once 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 then the manual reset will only be accepted after the system has generated a recovery event. Manual recovery is achieved by clicking **Manually Unfail** on the Dashboard. This tool is only be enabled if **Manual Unfail** is allowed and will have an effect if the condition that triggered the failure has not been resolved.

Example Deployment Configurations

This section provides examples of how the SSL Appliance can be deployed alongside other security appliances in order to protect the network against threats carried by SSL traffic.

In all the examples network links shown in red indicate links that are carrying decrypted SSL traffic.

Network port	Network interface which is either part of the “bump in the wire,” or is connected to a network tap device.
Device port	Network interface connected to the primary attached appliance which is handling inspected traffic from the SSL Appliance.
Copy port	Network interface which is connected to a secondary passive appliance, which receives a copy of the inspected traffic
Aggregation port	Network interface providing a connection to an additional network tap so the segment can receive traffic from more than one tap
Symmetric traffic	Packets for both directions of a network flow are seen on the same network interface on the SSL Appliance.
Asymmetric traffic	Packets for each direction of a network flow are seen on different network interfaces on the SSL Appliance.
Active-active	Packets on a given flow may be sent over either HA network link. From an SSL inspection perspective, this is equivalent to the asymmetric traffic scenario, in that packets belonging to a single flow might arrive on either one of two network interfaces.

Outbound Inspection

The next figure 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.

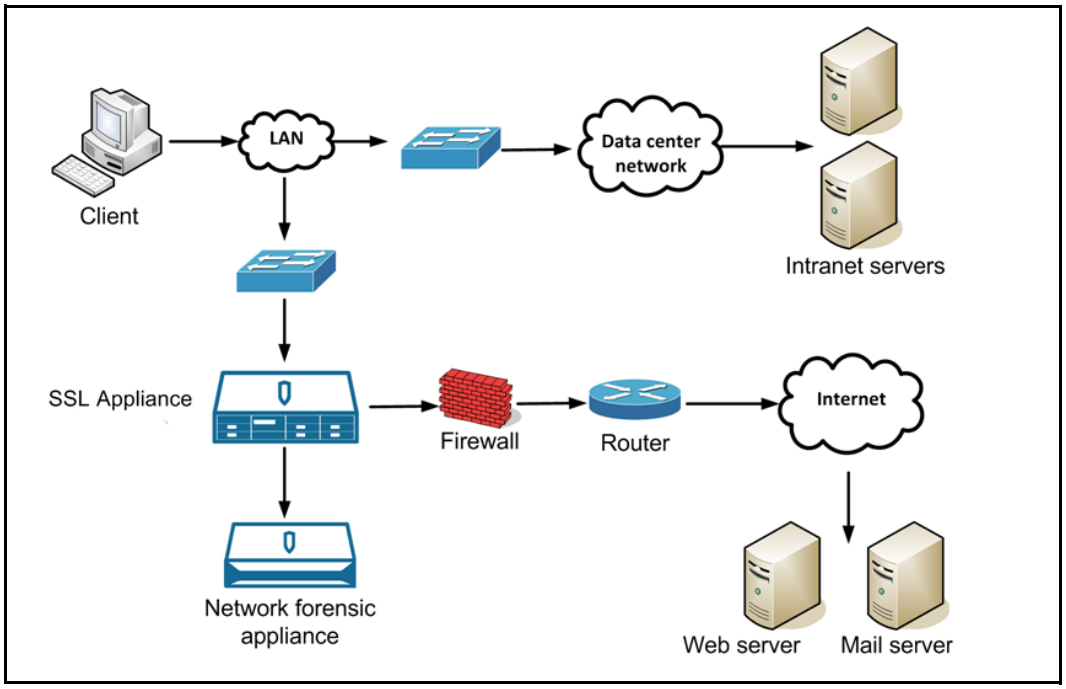

For this scenario the SSL Appliance is typically deployed adjacent to the firewall or router which leads to the Internet. The SSL Appliance needs to be deployed on the public side of the firewall if the firewall itself generates SSL encrypted traffic which needs to be inspected (for example, if the firewall also includes SSL VPN capabilities) or if the network topology requires deploying the SSL Appliance at that location (such as 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 would be inspected using certificate resign (see Certificate Resigning Method, page 3-4) as the SSL servers are not under the control of the enterprise deploying the appliance, so it is not possible to obtain copies of the server private key/certificate for these servers. The client systems in this deployment will need to trust the Certificate Authority used by the SSL Appliance to resign server certificates. The Symmetric Passive-Inline connection mode (Passive-Inline Mode, page 3-10) is used in this example.

Inbound Inspection

Inbound inspection is 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.

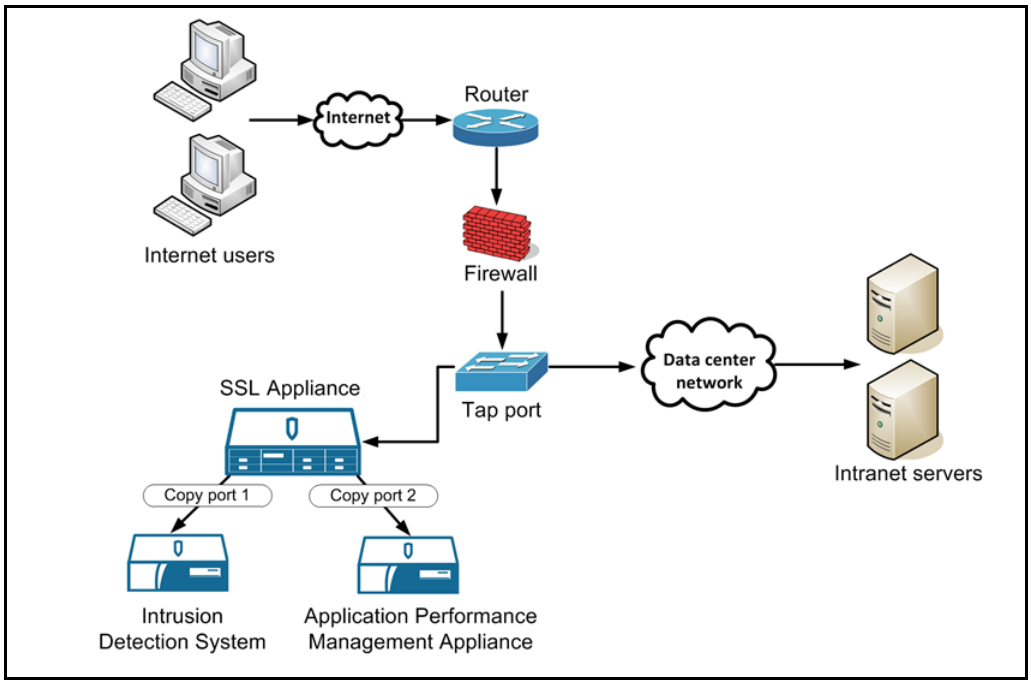

The private key and certificate for each of the Intranet servers are loaded into the appliance as it is using known server key mode to decrypt the traffic. Symmetric Passive-Tap connection mode (Passive-Tap Mode, page 3-8) is used in this example.

Inbound and Outbound Inspection

In an inbound and outbound deployment, both inbound and outbound traffic are inspected. The IPS in this deployment will be able to detect any threats in inbound sessions heading for the Intranet servers from users on the Internet and at the same time will be able to detect any inbound threats over sessions from users on the LAN to Internet servers. In addition the Network Forensic system will be able to detect and identify any files sent out as webmail attachments by internal users.In this example, the SSL

Appliance will be using both certificate resign and known server key mechanisms to decrypt traffic with the selection of which mode to use being determined by whether an SSL session is incoming or outbound.

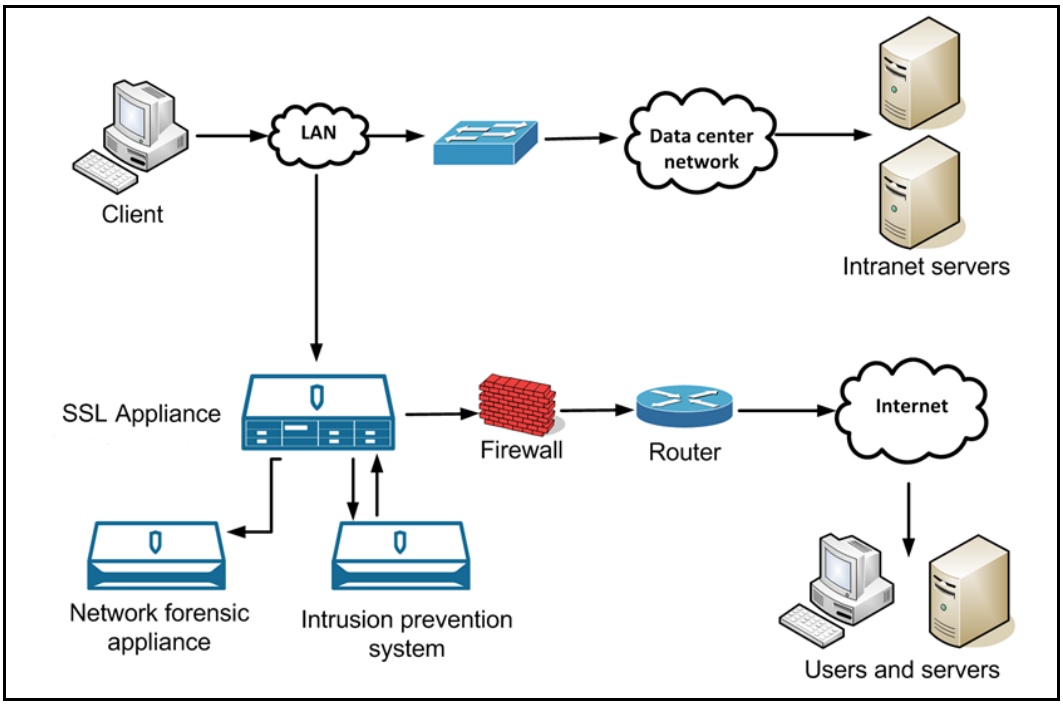

High Availability Deployment

Although an SSL Appliance segment has fail to wire capabilities provided by the Netmod to ensure connectivity, in most scenarios where hardware has failed or software is temporarily not available, some customers prefer to deploy multiple SSL Appliances, as this will ensure that in these scenarios traffic continues to be inspected.

Key to this deployment is having the SSL Appliance segment configured in HA mode with the software failure mode set to “Disable Interfaces,” and with link state mirroring enabled on the Ethernet switch devices.

Normally switch A1 and A2, SSL Appliance A it's attached security appliance(s) will be active. Should

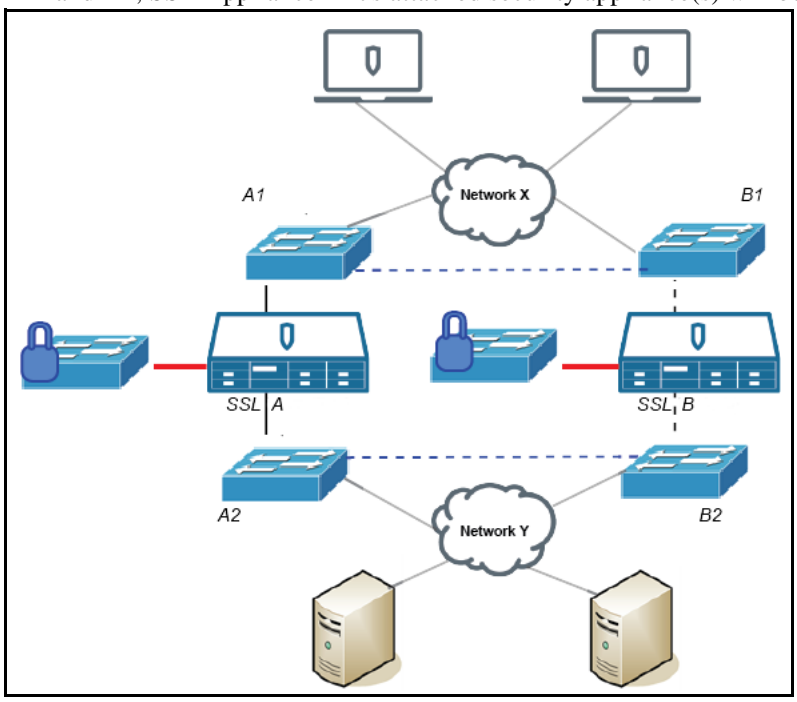

any of the links along that path fail, or should the SSL Appliance or it's attached security appliance or either of the Ethernet switches fail, the link down state will propagate, with standard mechanisms like the Spanning Tree Protocol or the Virtual Router Redundancy Protocol, ensuring that traffic is rerouted over the link between switches B1 and B2 that passes through SSL Appliance B (dashed line in the figure).

Availability can be further improved by including additional links between switch A1 and B1 and between switch A2 and B2 (shown as dashed lines in the figure). 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 might be combined into a single device, with A2 and B2 being similarly combined.

Contact Cisco support (tac@cisco.com) should you require more information with respect to High Availability deployment options.

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 (see User Interface Overview, page 6-1 for information on all features).

The front panel keypad and display can be used to configure the settings for the device and are also used during initial bootstrap mode and to unlock the master key during system start up.

A limited Command Line Diagnostic Interface (CLD: Command Line Diagnostics Interface, page 7-8) is also available. The *Cisco SSL Appliance Getting Started Guide* contains detailed information on using a console and the CLD to access the appliance.

Power On the Appliance

The very first time you plug the appliance in, it will start up automatically. After that, how an appliance behaves when power is restored after an outage varies with the model, as described in the following tables.

Table 5 SSL1500 Power On Behavior

State Before Power was Lost	State After Power is Restored
Booted and operational	Booted and operational
Powered down using the WebUI or CLD	Powered down
Powered down using the power button	Booted and operational

Table 6 SSL2000 and SSL8200 Power On Behavior

State Before Power was Lost	State After Power is Restored
Booted and operational	Booted and operational
Powered down using the WebUI or CLD	Powered down
Powered down using the power button	Powered down

If the appliance is powered down when power is restored, use the physical power button to power it back on.

Bootup Behavior

Note The SSL Appliance is factory configured to use DHCP to acquire an IP address for the management Ethernet.

You will see several start up messages on the LCD screen before the appliance boots up. See the *Cisco Getting Started Guide* for detailed guidance on initial setup. The front panel LCD with the default screen that is displayed in normal operation once the bootstrap phase is complete. Here are two sample screens.

Once the appliance is through the initial bootstrap phase (including setting up the management interface IP address), you can access the WebUI with a browser. You will want to complete the following tasks, setting the date and time, verifying or editing the management network settings, and setting up one or more management users.

Configuring System Date/Time and Timezone

To configure the system date and time use the **Date/Time** option on the **(Platform Management)** menu (under the system name).

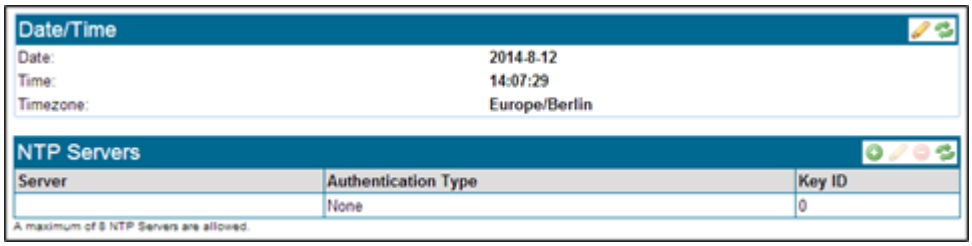

If NTP is enabled, the **Date** and **Time** fields will be disabled as these values are being set by the Network Time Protocol (NTP). Click the Add (pencil) icon to edit these settings. Configure the settings then click OK to save the settings. The screen will refresh.

See Date/Time, page 6-61 for details on setting the date and time, and for using NTP servers.

Note If you have changed the date, time, NTP, or timezone, you must select Apply at the “Platform Config Changes” message which displays at the bottom of the screen.

Configure Management Network Settings

To configure system settings, use the **(Platform Management) > Management Network** menu (see Management Network, page 6-49 for more information on all of the settings).

By default, the appliance management network requests DHCP addresses for both IPv4 and IPv6. The IP address displays on the LCD screen on the front of the appliance. See IPv6 Settings Panel, page 6-51 for more information on the IPv6 options, and see the *Getting Started Guide* for your appliance for further details.

The SSL Appliance supports simultaneous access by both IPv4 and IPv6. If both IPv4 and IPv6 are disabled (**Mode > Disable**), access is only possible through the Command Line Diagnostic Interface (see Command Line Diagnostics Interface, page 7-8) if this can be accessed directly from the console, or via the serial console (see the *Getting Started Guide* for your platform).

The next figure shows an active IP configuration window on top of the **Management Network** window. In this example, the system is configured to automatically obtain IPv4 settings using DHCP. The **IP Address/Netmask, and Default Gateway** fields can't be edited.

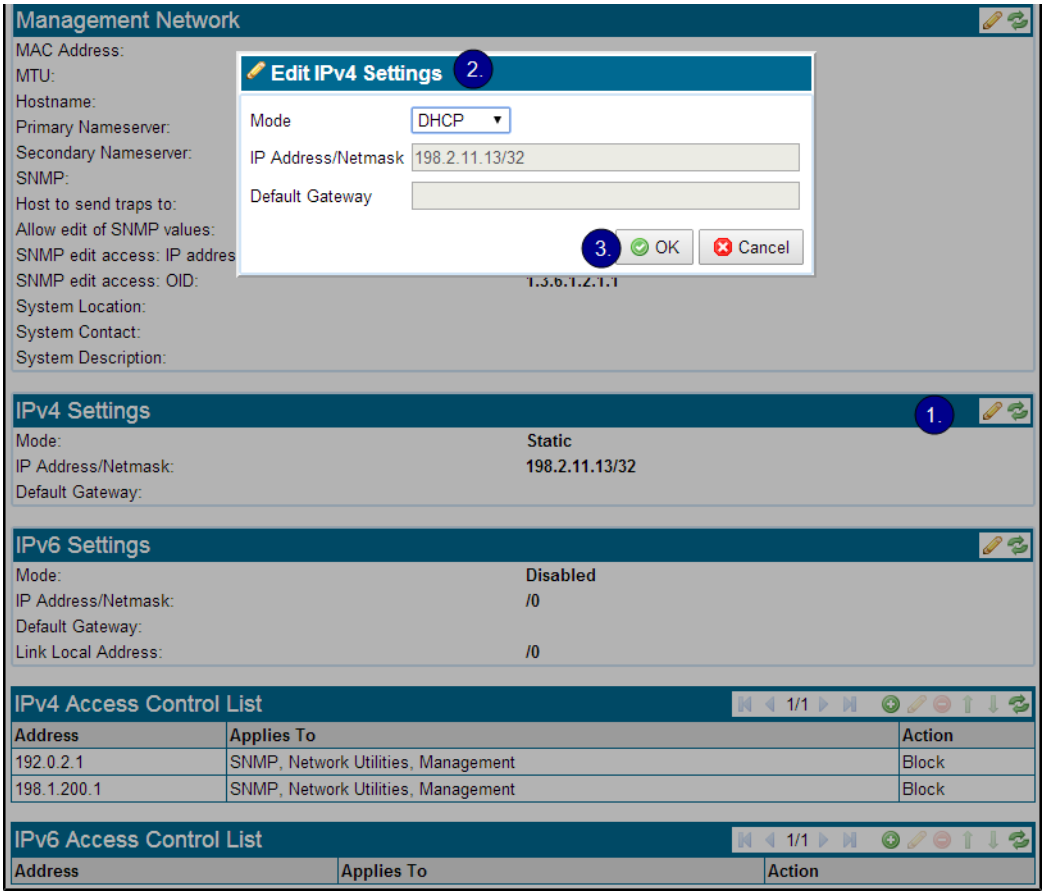

Note Changes to any IP network settings require an appliance restart.

Configure a Static IP Address

Step 1 Click Edit on the required IP (**IPv4** or **IPv6**) **Settings** header; the **Edit** window opens.

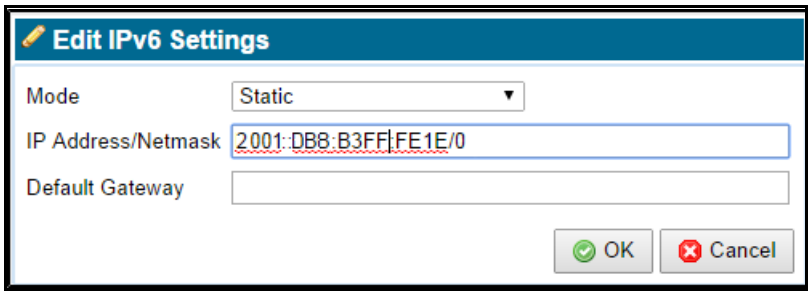

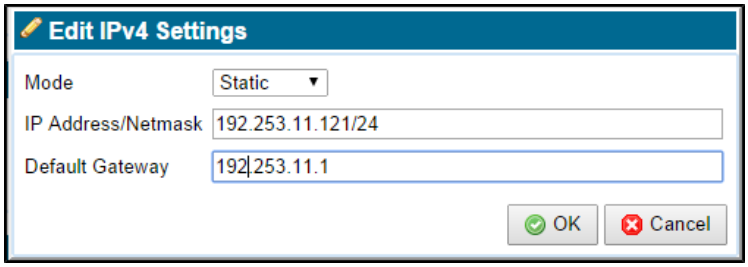

Step 2 For Mode, select **Static**.

Step 3 Enter the required data, as shown in the table.

The **IPv6 Link Local** address shown in the **IPv6 Settings** panel is derived automatically and presented for IPv6 settings; you may not edit it.

IPv4	IPv6
IP Address/Netmask	IP Address/Netmask
Default Gateway	Default Gateway
Note: Use the IP address/mask bits (CIDR) format to enter the IP address and netmask. Example: 192.0.2.0/24. See IP Address Format Notes, page 6-50 for more details on entering IP addresses.	

Step 4 Click OK. The **Edit** window closes.

Step 5 Click **Apply** to save and apply the changes. The changes to the network settings will only take place once the reboot has occurred.

Note After you click Apply, you must reboot the appliance. Click Reboot on the apply changes message.

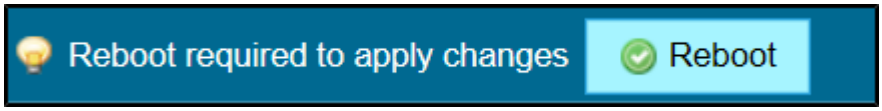

Use **Access Control Lists** to control incoming connections on the management interface. You can create blacklists (never allow a connection) and whitelists (always allow the connection). See IPv4 and IPv6 Access Control List Panels, page 6-51 for detailed configuration information.

Configure Management Users

Create new user accounts on the system using the **Users** option on the **Platform** menu. Click on the + icon to add a new user to the system.

The next figure shows the **Add User Management** window with the details required to add a user. The Roles section lets you assign one or more roles to the user being created.

Note To complete the bootstrap phase, you must create users with the **Manage Appliance** and **Manage PKI** roles. These roles can be assigned to the same user or different users.

To assign more than one role click the first role, which will highlight the role, then hold down the CTRL key (Command key, for Mac users) and click on a second role which will also be highlighted. Repeat this process until all the roles you wish the new user to have are highlighted and then click **Save**.

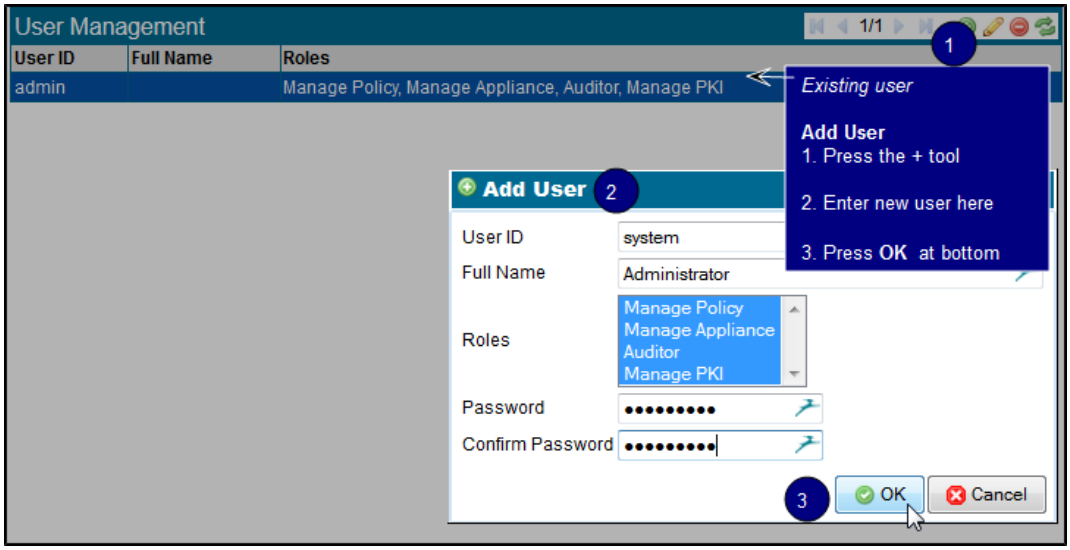

Click OK to create and add the new user to the system.

See User Management, page 6-76 for information about the privileges of different roles.

Users 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 the menu at the top right of the screen under the user name.

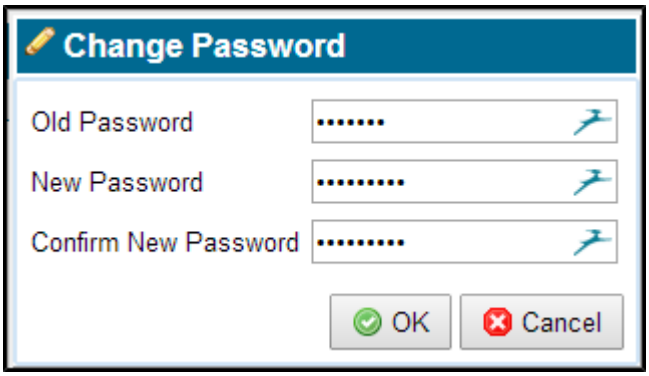

System Status

To view the overall status of the appliance, click on the **Monitor > Dashboard** menu option (Monitor the Dashboard for Current Status, page 6-6). Status details shown here feed into the summary status indicators for **System, Load, Network**, and **License** that appear in the footer at the bottom of the display. The Appliance Uptime is displayed just below the system graphic at the top of the window.


Monitor
Policies
PKI
Appliance Uptime: 3 days, 21:34:02
Segments Status
Segment ID
Main Interfaces
Copy Interfaces
Interfaces Down
Main Mode
Failures
Network Interfaces
Port | Type | Link State | RX Packets/Bytes | TX Packets/Bytes | RX Drops
1 | 1G | Down | 0/0 | 0/0 | 0
2 | 1G | Down | 0/0 | 0/0 | 0
3 | 1G | Down | 0/0 | 0/0 | 0
4 | 1G | Down | 0/0 | 0/0 | 0
5 | 10G | 10G | 0/0 | 0/0 | 0
6 | 10G | 10G | 0/0 | 0/0 | 0
7 | 1G | 1G | 0/0 | 0/0 | 0
8 | 1G | 1G | 0/0 | 0/0 | 0
9 | 1G | 1G | 0/0 | 0/0 | 0
10 | 1G | 1G | 0/0 | 0/0 | 0
CPU Load %
cpu | cpu0 | cpu1 | cpu2 | cpu3 | cpu4 | cpu5 | cpu6 | cpu7 | cpu8 | cpu9 | cpu10 | cpu11 | cpu12 | cpu13 | cpu14 | cpu15 | cpu16 | cpu17 | cpu18 | cpu19
1.3 | 11.9 | 0 | 17.3 | 0 | 1 | 2 | 0 | 0 | 1 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0
Fan Speed (RPM)
Platform Config Changes
Apply
Cancel
2015-11-11 02:09:04 Blue Coat Systems, Inc. SV3800 3.9.2.1-18
System: Load: Network: License:

Install a Local CA for Certificate Resign

Before the SSL Appliance can be used to inspect traffic using Certificate resign mechanisms it must have at least one CA certificate and private key installed which can be used to do the resigning. 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 resign installed then it is possible to use different CAs to resign different SSL sessions by choosing the appropriate CA in the policy configuration.

Management of local Certificate Resigning Authorities is done with the **Local Resigning Certificate Authorities** option on the **PKI** menu. The SSL Appliance might also work with an HSM appliance; see Work with a SafeNet Java HSM, page 5-1.

If the SSL Appliance is operating in an environment where SSL server certificates signed by the CA using an EC key are present, you must create or load one or more resigning CAs which use EC keys. When creating a self-signed CA on the appliance, you can specify if the CA should use RSA or EC keys. The type of key being used by an resigning CA is shown on the WebUI.

Local Resigning CA Specific Tools

In addition to basic list edits, when you edit a Resigning Certificate Authority, you have the option to add a **CRL URL**; this URL location is inserted into any re-signed server certificates signed by the resigning CA.

Generate a new Resigning Certificate Authority:

Add a Resigning Certificate Authority by importing an existing CA and key

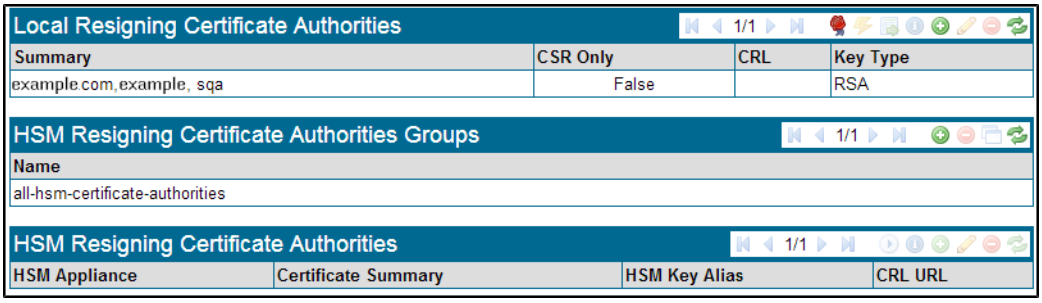

The following subsections consider each of these ways of adding a Resigning Certificate Authority.

Create a Local CA

Go to **PKI > Resigning Certificate Authorities**. In the **Local Resigning Certificate Authorities** panel, click the rose icon to generate a CA, bringing up the **Generate Certificate** window. Enter the basic data required in a CA, including the key size and validity period. Once the data is input, choose to:

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

If you select **Generate self-signed CA**, there are no further steps. The CA is generated and added to the set of resigning certificate authorities in the system. As 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. SeePKI Management, page 6-42 for details on how to do this. Click Apply and the certificate is saved and installed, and an entry in the **Local Resigning Certificate Authorities** table displays, with an indication that no CSR has been generated for this certificate.

If you select **Generate a CSR**, a PEM format CSR is generated. It needs to be sent to the Certificate Authority that is going to sign it.

Copy the text in the **Certificate Signing Request PEM** field into a file. The file then must be communicated to the CA that will sign the final Resigning Certificate Authority certificate. When you click OK, the certificate details are saved, and an entry in the resigning certificate authorities table displays with an indication that a CSR has been generated for this certificate. At this point the certificate is not installed in the system, as the signed resigning CA has not been received back from the CA that is signing it. When an entry in the table shows CSR True, the install a certificate icon is active. When used, you will prompted to provide the signed CA so it can be installed in the system.

Caution 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 then the CSR should be presented to the private root CA and the private root CA should sign this to create a private Intermediate CA which can then be loaded onto the SSL Appliance and which the client machines will trust as it is signed by the private root CA that they already trust.

Note 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 **Local Resigning Certificate Authority** screen displays in the list, as it does in the previous figure. At this point the CA cannot be used, as the signed certificate from the CA that the CSR was sent to has not been loaded. Once the signed certificate is available it can be loaded by selecting the entry in the **Local Resigning Certificate Authorities** window and clicking the icon, allowing the signed certificate to be imported into the system. See the next section.

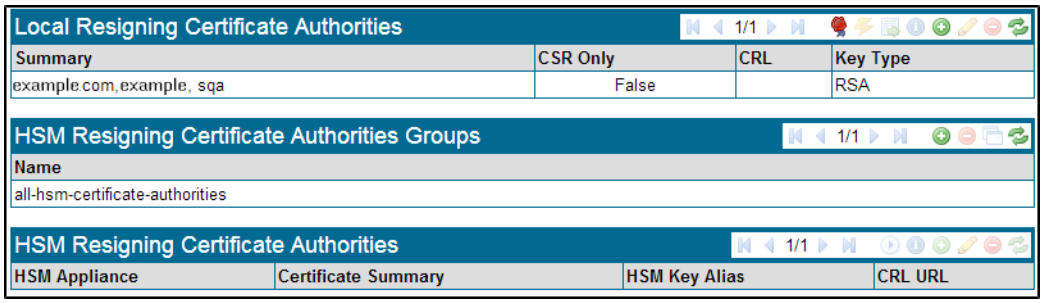

Import a CA

If you already have a CA that you want to use as an Resigning Certificate Authority in the SSL Appliance, you can import it and install it in the system. You will need both the CA certificate and the private key for the CA in order to install it on the system.

Click Add in the **Local Resigning Certificate Authorities** to bring up the **Add Local Certificate Authority** window. You can either browse to and upload the files containing the certificate and private key, or paste the certificate and private key in directly.

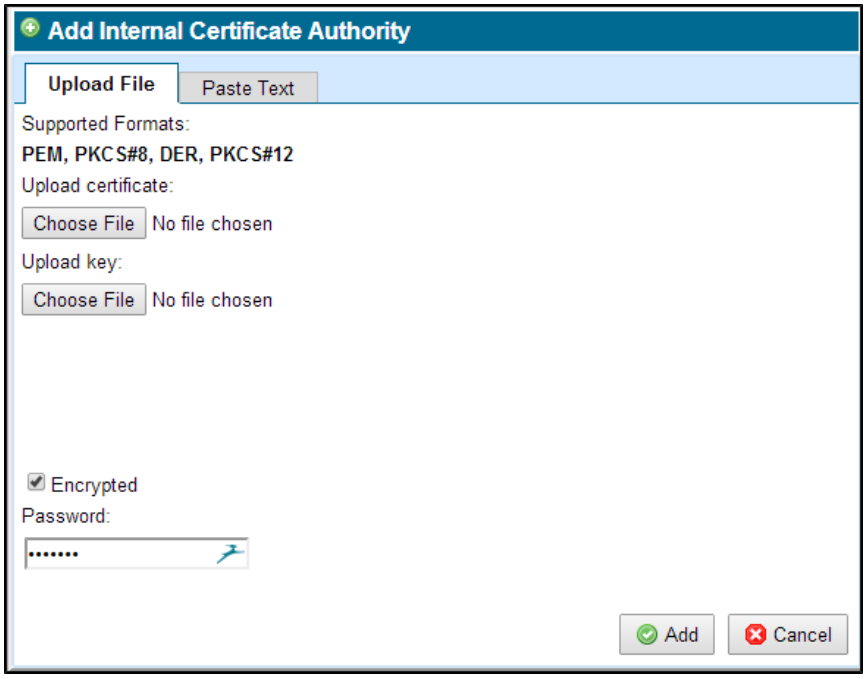

If the certificate and key being imported have been encrypted and protected with a password, you must check **Encrypted** and type the password in the **Password** field.

Import Known Server Keys

In order 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 are loaded into the SSL Appliance. Use the **Known Certificates and Keys** option on the **PKI** menu to import new certificates and keys. Known server certificates and keys are imported into the **all-known-certificates-with-keys** list, and can then be copied to custom lists if required.

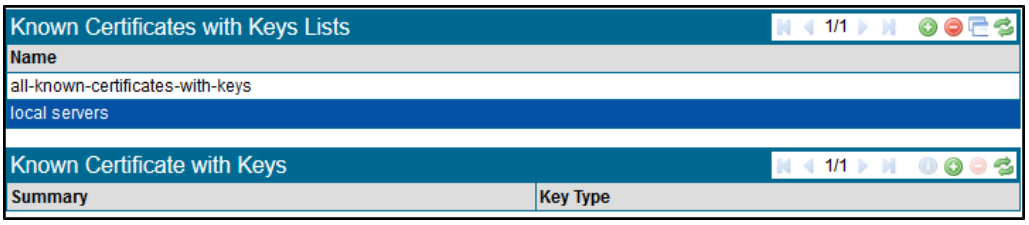

There are two panels, one to choose the list that is to be operated on and the other to manipulate the contents of that list. Initially there will only be one list, called **all-known-certificates-with-keys;** it will have no certificates in it.

In order to import the first known server key and certificate, click the all-known-certificates-with-keys entry in the **Known Certificates with Keys List** window, then click **Add**.

The **Add Known Certificate with Key** window displays. 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 displays, including a **View Details** button. You will see that the key displays as a row in the **Known Certificates with Keys** panel.

Click **Apply** to save the imported certificates and keys to the secure store.

PKI Management, page 6-42 explains how to create custom lists of Certificates and Keys in more detail.

Example Passive-Tap Mode Inspection

The following example shows the steps for configuring the SSL Appliance to inspect traffic that is destined for a server that you can obtain a copy of the private key and certificate from. In this example the SSL Appliance is deployed in passive-tap mode with an additional copy port, as described in Passive-Tap Mode, page 3-8. The known server certificates and keys used in this example are shown in the figure.

The steps involved are:

- Load the server key/certificate into the SSL Appliance (see Import Known Server Keys, page 4-12)
- Create a ruleset that contains a rule to inspect traffic to the server
- Create a segment for passive-tap operation
- Activate the segment to start inspection

In this example the certificate and key for ***bluecoat.com*** is used to allow inspection of traffic going to that server. As this certificate/key is already loaded into the system, we can proceed to the next step, which is to create a ruleset that contains a rule specifying that traffic to bluecoat.com should be inspected.

This is a two step process, first creating the ruleset to hold the rule, then defining the rule itself. The next figure shows the screen while adding a new ruleset called passive-tap-example. After clicking OK the new entry will appear as a row in the **Rulesets** grid, and is available for use. At the bottom of the screen is a **Policy Changes** notification area with options to **Apply** or **Cancel** the change. Click **Apply** to complete the process, and to save the ruleset to disk.

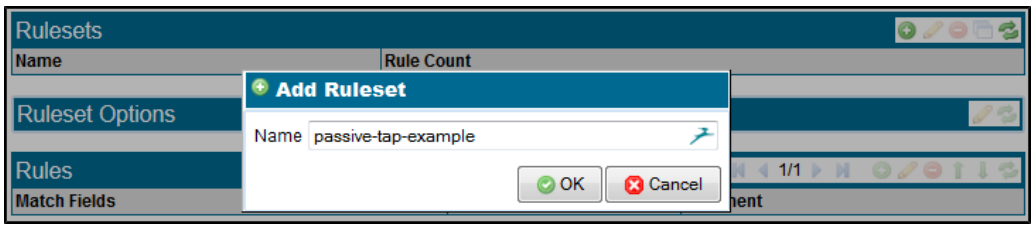

Next, click the **passive-tap-example** row to select it. This will display the **Ruleset Options** for this ruleset. In this example the default settings are fine, and are explained below:

- No **Resigning Certificate Authority**, as we are not doing certificate resigning
- All **External Certificate Authorities** and CRLs are used when checking an SSL session
- There are no trusted certificate 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 then the default setting would be to use All Trusted Certificates.
- Any SSL sessions that don't match a rule in this ruleset will be cut through to the attached security appliance without being decrypted

Clicking **Add** in the **Rules** grid section opens the **Insert Rule** pane. Select **Cut Through** on the **Action** drop down menu to configure the valid options for this rule.

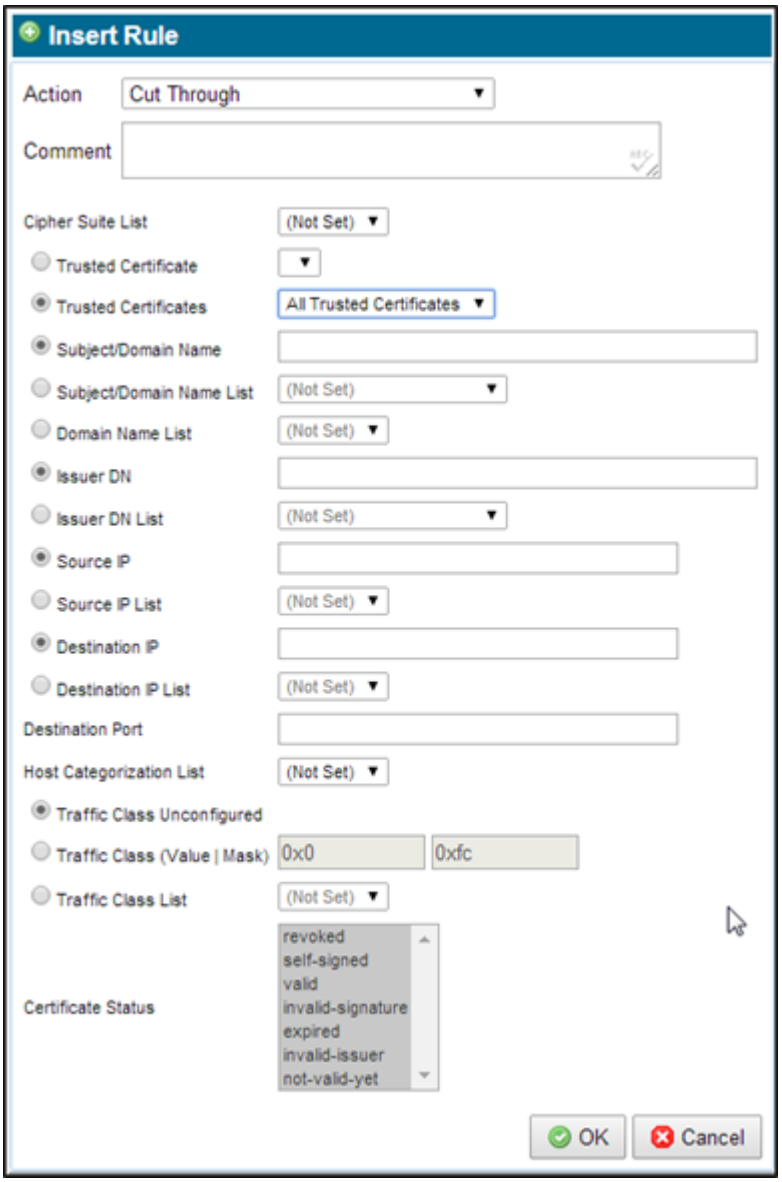

In this example the rule only applies to a single server for which the certificate and key are known, so the **Known Certificate with Key** option is checked, and the system for which we loaded the key is selected from the drop down menu. Apart from adding a comment to the Comment field, no other options are used in this rule, so click **OK** to create the rule. At the bottom of the screen is a **Policy Changes** notification area. Click **Apply** to complete the process and save the rule to disk.

The final part of the process is to create a segment, configure it to use the ruleset just created and then to activate it. To create a Segment go to the **Policies > Segments** menu option; you will see the **Segments** information.

The next figures show the segment screen when no segments currently exist on the system. The graphic at the top of the screen identifies the appliance model. The ports that show green on the graphic indicate that the links on these ports are up.

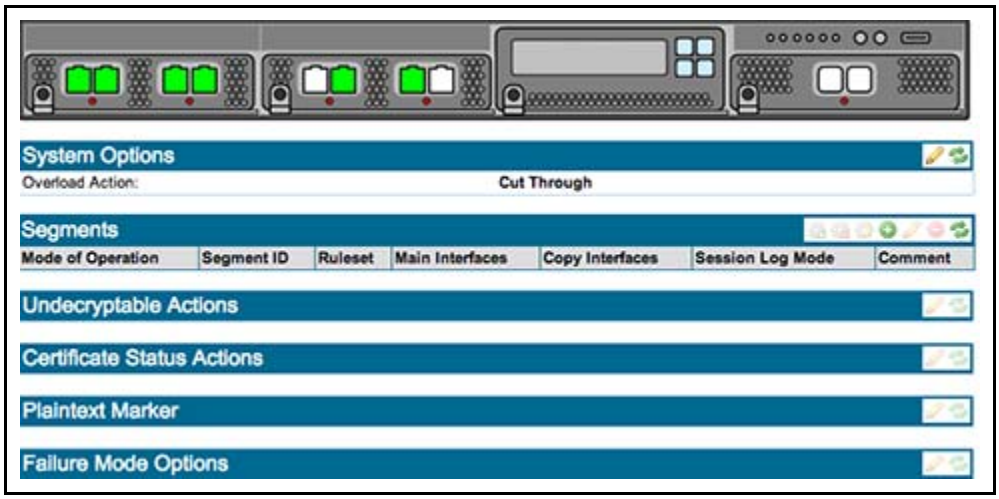

To create a new segment, click **Add** in the **Segments** table.

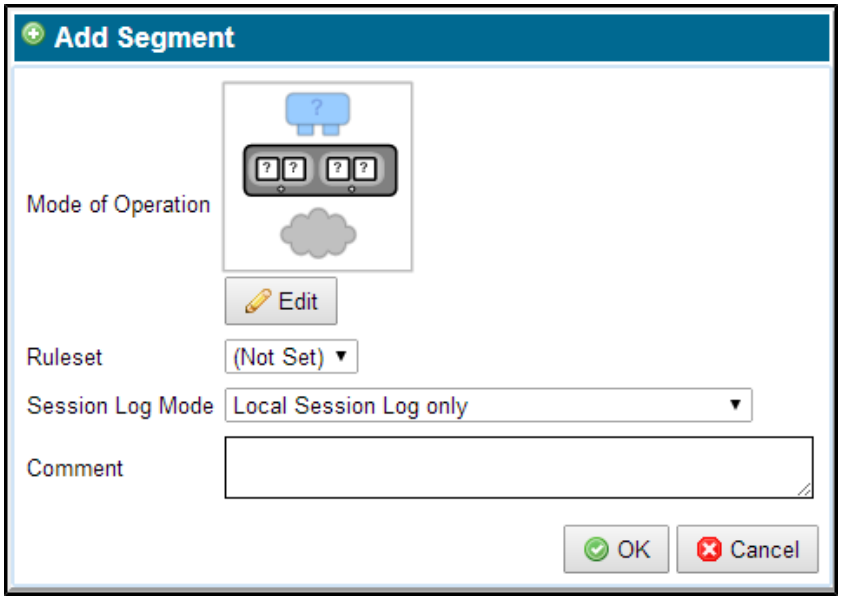

To select the **Mode of Operation** click **Edit**, then choose the required mode from the **Select Mode of Operation** menu. Choose the **Ruleset** from the drop down menu.

In the **Select Mode of Operation** window, choose the mode of operation for a segment.

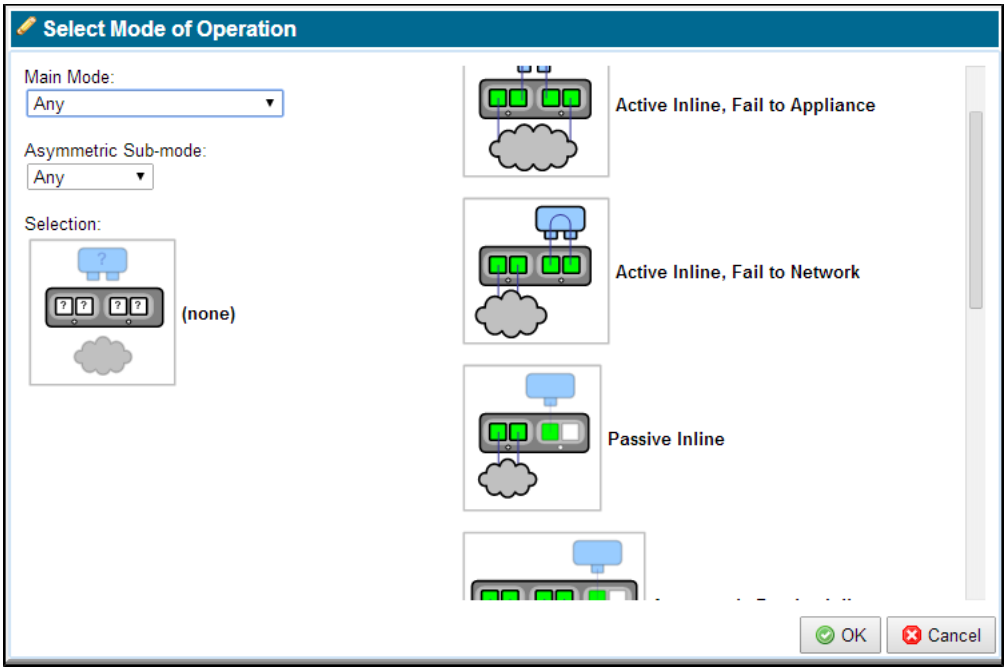

The **Select Mode of Operation** window displays the different operating modes as images. Narrow the set of operating modes using the **Main Mode** drop down menu. Choose Passive Tap for example; this will reduce the number of options displayed in the **Mode of Operations** area. You can use the **Asymmetric Sub-mode** drop down menu to further narrow the number of modes of operation. Click the image of the desired operating mode to select it, and click **Save** to set this as the mode of operation for the segment.

Select the **Ruleset** as required.

Logs can be saved locally, or you can send errors or session logs to remote servers, at the **Session Log Mode** field. Make sure to follow up with the **Remote Logging** menu item (Remote Logging, page 6-61) to actually transmit the logs remotely.

If you want to translate VLAN tags between ports, select **VLAN Translation**, and see Translate VLAN IDs with VLAN Mappings, page 6-28.

The next figure shows the completed segment details before they are saved. In this example, the local session log has been selected. The graphic in the input window indicates that this segment will make use of two ports on the system. The actual port numbers to be used are not known at this point; they are determined when the segment is actually activated.

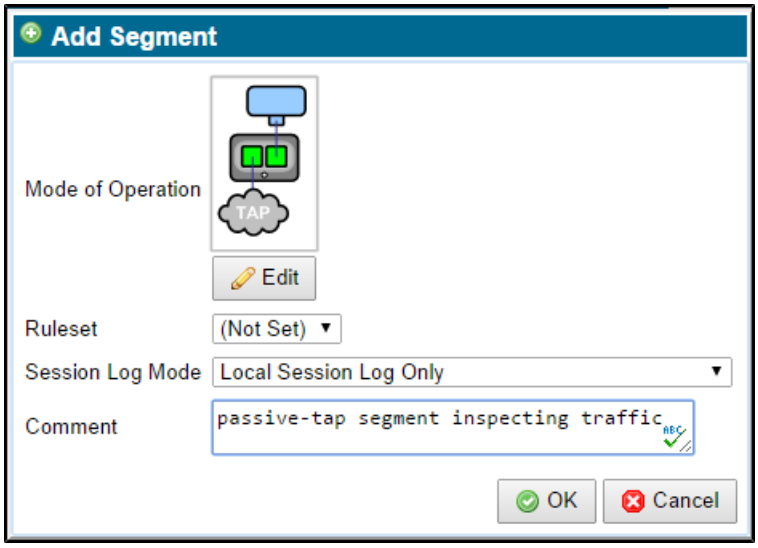

Click **OK** in the **Add Segment** window to create the segment. At the bottom of the **Segments** screen is a **Policy Changes** area, with options to **Apply** or **Cancel** the change. Click **Apply** to complete the process and save the rule to disk.

Once created, you can see the segment in the **Segments** table, and can select it by clicking on it. There are several panels below the **Segment** table, each of which allow different types of actions to be configured for the selected segment. These are explained next. To change any of the settings in the panels, click **Edit** for that panel.

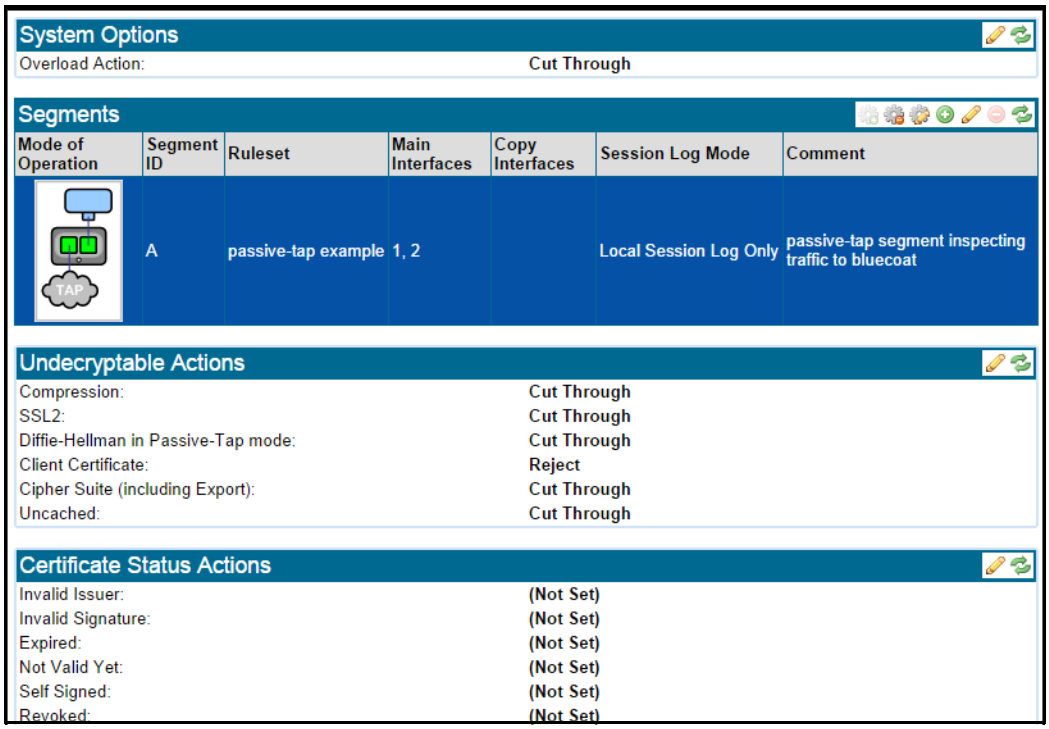

The **Undecryptable Actions** panel gives you control over what will happen to an SSL session that cannot be decrypted by the SSL Appliance. Different actions can be configured depending on the reason why decryption is not possible. In this example, the action is to cut through the session except in the case where client certificates are used when the SSL session will be rejected.

The **Certificate Status Actions** panel gives you control over what will happen 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. Use Status Override Order to configure which Certificate Status actions have priority, those configured for the segment, or those configured in a rule in the ruleset in use 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, so the override setting and segment default actions have no effect.

The **Plaintext Marker** panel lets you control how the generated flow with the decrypted payload is marked, of if it is marked at all. The options are to have these flows be marked with:

- VLAN tag; the VLAN ID used is configurable
- Modified source MAC address
- No marking

As 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 as generated flows must be marked in order that the SSL Appliance can identify them when they are sent back to it by the attached security appliance.

In this example, the generated flows will be sent out with no marking.

Notice that the **Interface** columns in the **Segment** do not shows interface numbers; these are allocated when the segment is activated. Click **Activate** for the segment to activate it, which is in the tool block at the top right of the segment panel, then click **Apply**.

During the activation process a series of screens appear for you to select the ports to use for the segment, and to select any copy ports and the modes that the copy ports will operate in. The initial screen indicates which interfaces on the device are available for use and which are already in use by other segments. For information on configuring VLAN translation, see VLAN Notes, page 6-29.

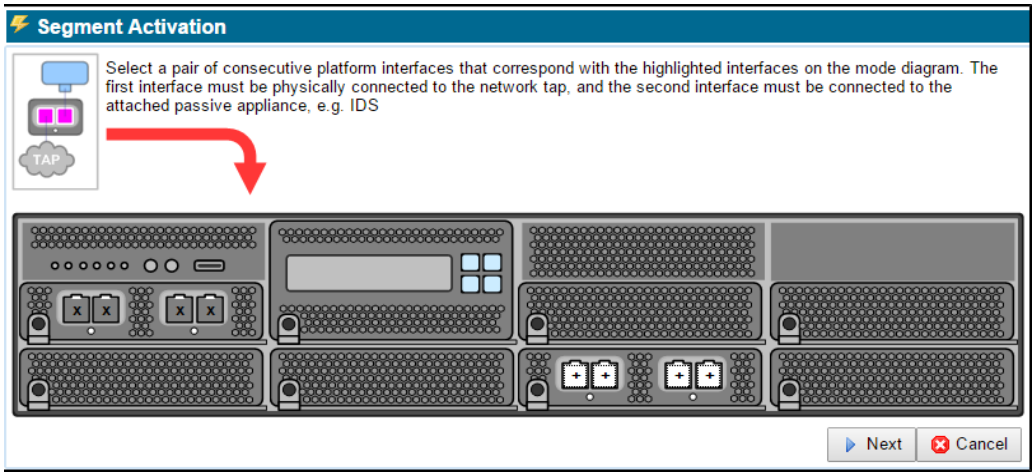

You can see next that ports 5 and 6 have been selected as the two primary ports for this segment. Click **Next** to move on to the next step in the process.

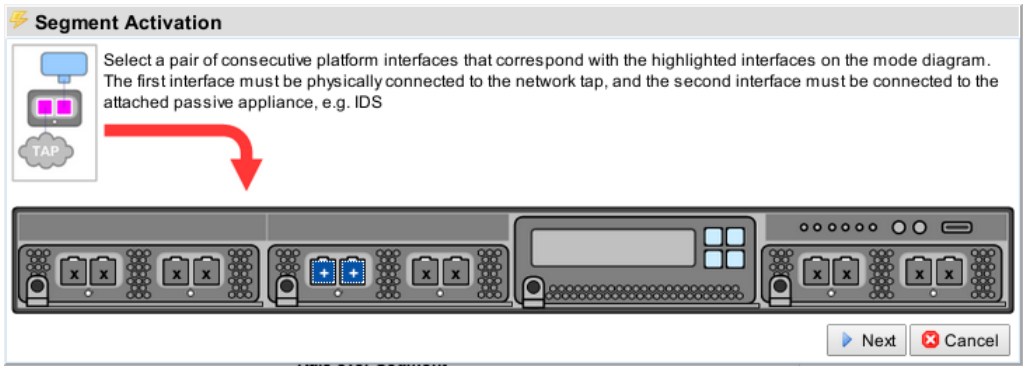

One or two mirror ports can be configured for this passive-tap segment, indicated by the images in the box at top left. One mirror port has been selected in this case. If two mirror ports had been selected then the options allowing selection of per-direction copy or load balancing would be active allowing selection of these capabilities if required. Click **Next** then **Apply** to finish the activation process.

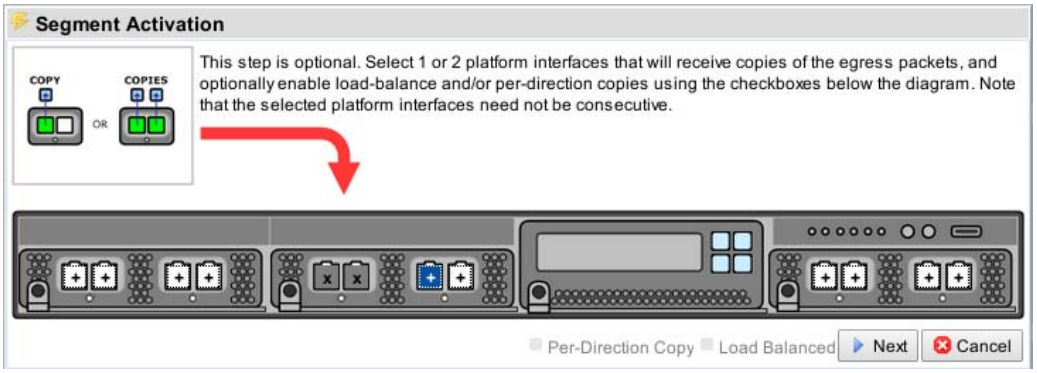

Once the segment is active, the **Segments** screen will show an entry for the new segment, and the graphic at the top of the dashboard will indicate the ports being used by the segment, seen next. In this example the segment is identified as Segment A, and is using three of the active ports, all of which show the letter A.

The green background indicates that this segment is activated. If there is SSL traffic to the server the **SSL Session Log** and **SSL Statistics** screens should show this. See SSL Session Log, page 6-10 for details on the session log and other monitoring tools.

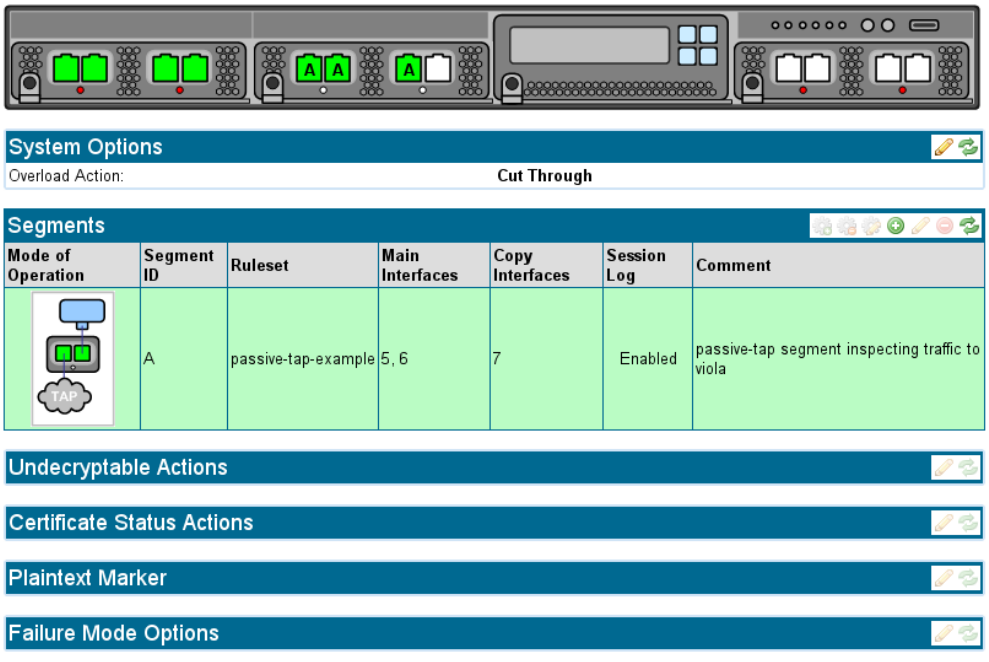

Example Passive-Inline Mode Inspection

The following example shows the steps for configuring the SSL Appliance to inspect traffic that is destined for a number of SSL servers that you cannot obtain a copy of the private key and certificate for. In this example the SSL Appliance is deployed in passive-inline mode as described in Active-Inline Mode, page 3-11. This example illustrates the use of certificate resign to inspect traffic and also 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 steps involved are:

- Create or load an resigning CA certificate and key into the SSL Appliance (see Install a Local CA for Certificate Resign, page 4-8)
- Create a ruleset that contains rules to inspect traffic going to specific destinations
- Create a list of destinations for use by a single rule
- Create a segment for passive-inline operation
- Activate the segment to start inspection

The next figure shows the edit options screen for a ruleset called passive-Inline-example that has already been added to the rulesets on the system. The resigning CA created above is selected as the default Resigning Certificate Authority.

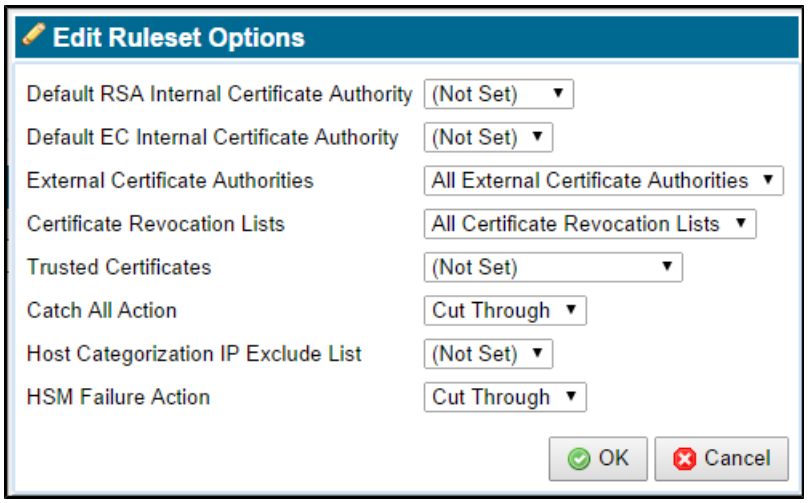

Before adding any rules to this ruleset we will create a list of Domain Names (DN) that will allow a single rule to apply to SSL sessions to multiple destinations.

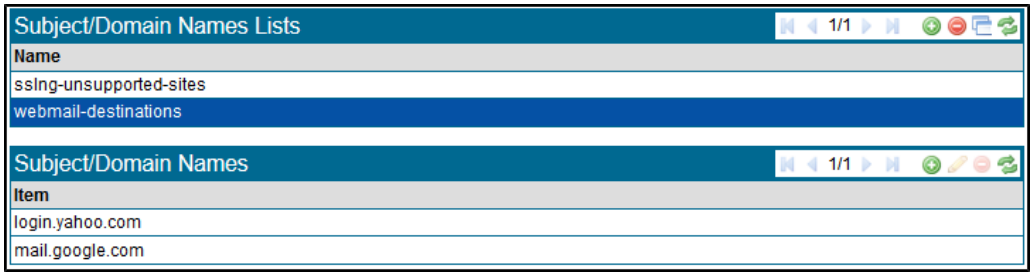

To create the list, click **Add** in the **Subject/Domain Names List** area, then give the new list the name "webmail destinations." After creation, select the empty list in the **Subject/Domain Names List** area, then click **Add** in the **Subject/Domain Names List** area, so a name is added to the list. Two **Subject/Domain Names** have been added to the list. At the bottom of the screen is a **Policy Changes** notification, with options to **Apply** or **Cancel** the change. Click **Apply** to complete the process and to save the new list to disk.

Now that the list exists, go back to the ruleset and add a rule to use this list. Verify that the radio button beside **Subject DN List** is selected, and select **webmail destinations** from the drop down menu.

In this example, the Destination Port could be set to 443. The effect of this rule will be to inspect any traffic going to a server that has a DN which is in the **webmail destinations** list and 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 then this rule would not be triggered.

Note In this example the entries added to the list are all Domain Names, and were simply typed into the **Add to List** window. You can include other elements of the X.509 certificate in a list by specifying what the item is when it is added. If the type of item being added is not specified, it is assumed to be a Common Name. More details on how to include other elements of the X.509 certificate in a list are given later in this document.

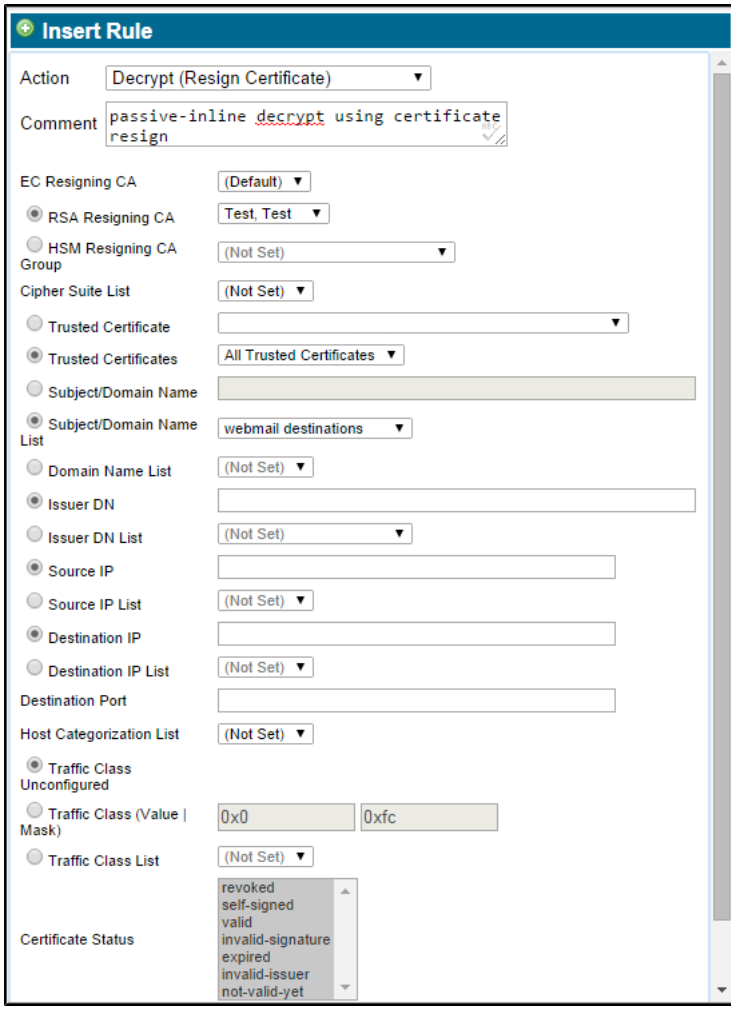

Having created the rule, click **OK**. As the default action for this ruleset is “cut-through”. Any SSL traffic which does not match the rule will be cut through and will not be inspected. If you wanted to prevent traffic to a specific SSL site then you could add another rule to the ruleset that matched on the specific Domain Name for that site, and had an **`Action`** to drop the traffic.

The next figure shows how the ruleset displays after a second rule has been added that will prevent any SSL traffic going to www.bluecoat.com.

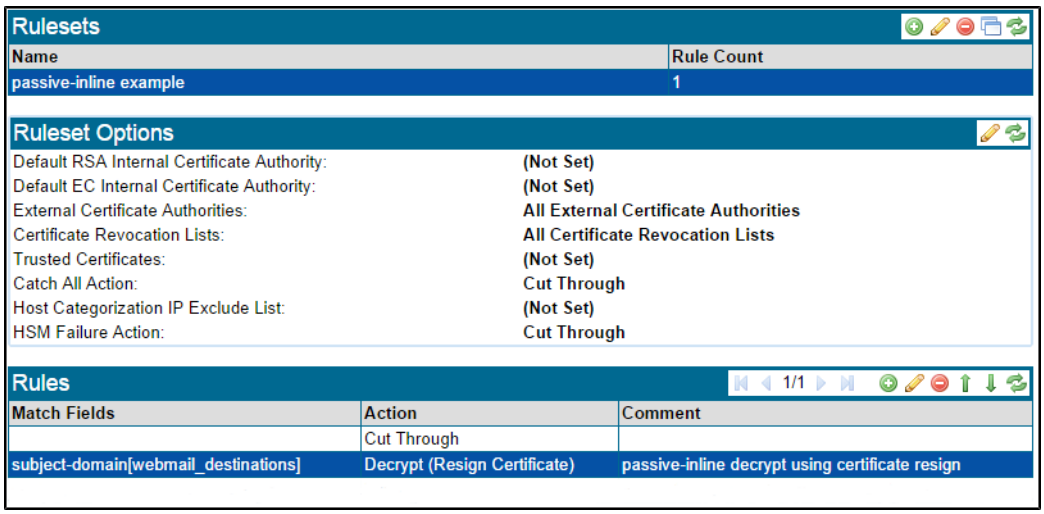

Having created the second rule, click **Apply** at the bottom of the screen. You will be able to see that the rules are now part of the ruleset.

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

To create a segment, go to the **Policies** menu and select **Segments**, then click **Add** in the **Segments** panel, and follow the same process as in the earlier example, but choosing a Passive-Inline segment type. Click **Apply** at the bottom of the screen to complete the process and to save the CA to disk. The next figure shows the segment after it has been completed and saved. Notice that:

- The ruleset created above is configured as the ruleset to be used for this segment.
- The local session log has been turned on for this segment
- The segment ID is B

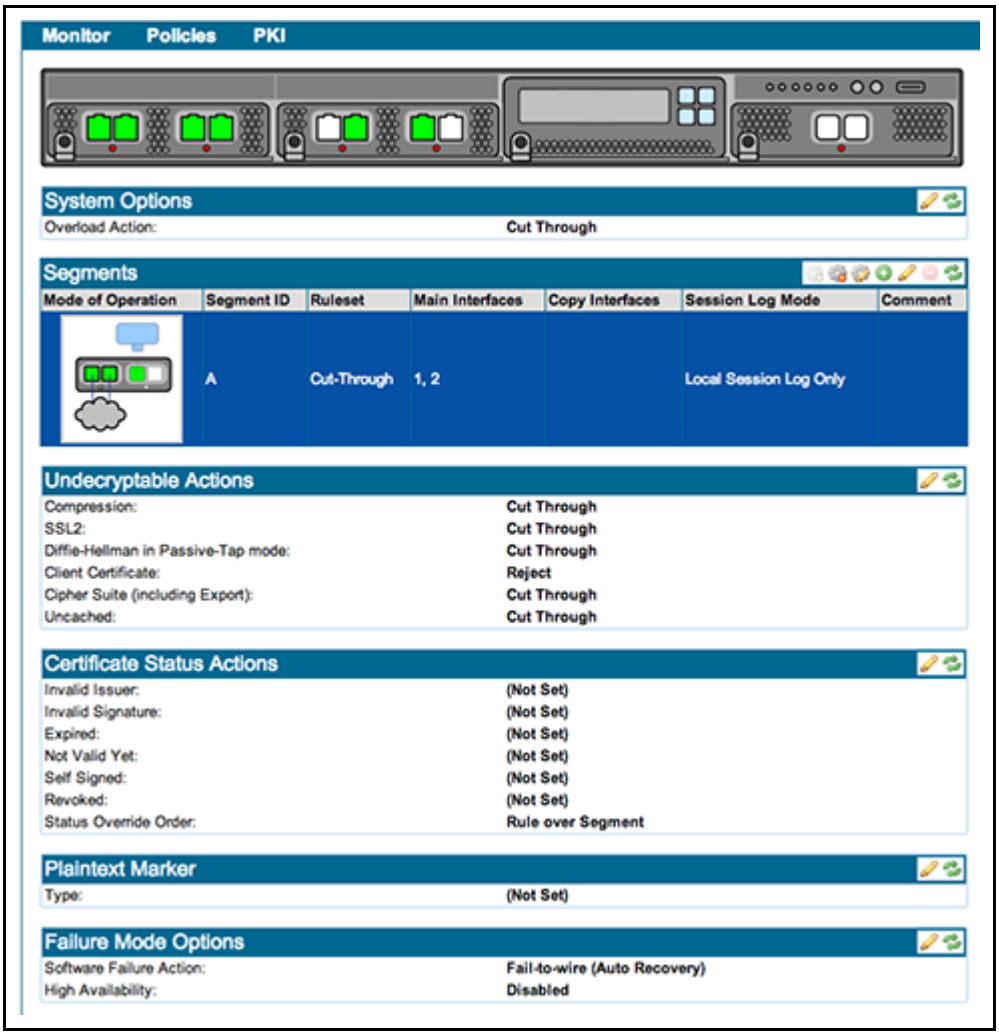

The final figure shows the segment status once it is active. The interface numbers indicate how the device should be wired up to the network. In this example:

- Interfaces 9 and 10 connect to the network making the SSL Appliance a bump-in-the-wire
- Interface 11 connects to the attached passive security appliance

The green background indicates that the segment is active. If there is SSL traffic to the server then the SSL **Session Log** and **SSL Statistics** screens should show this. See Section SSL Session Log, page 6-10 for details on the session log and other monitoring tools. The details for the passive-inline segment configured in an earlier example (segment A) are shown on the next figure.

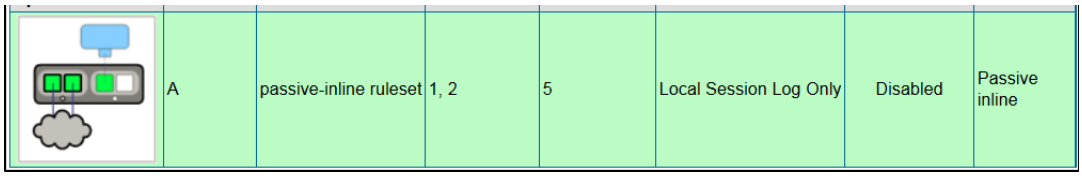

Example Active-Inline Mode Inspection

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, page 3-11. This example illustrates the use of both certificate resign 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.

If you are using VLAN for your setup, see About VLAN Configurations, page 6-28.

The steps involved are:

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

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

- Create a list of known server key/certificates
- Create a ruleset that includes both known server key inspection and certificate resign inspection
- Create an inline-active segment

These steps are shown next

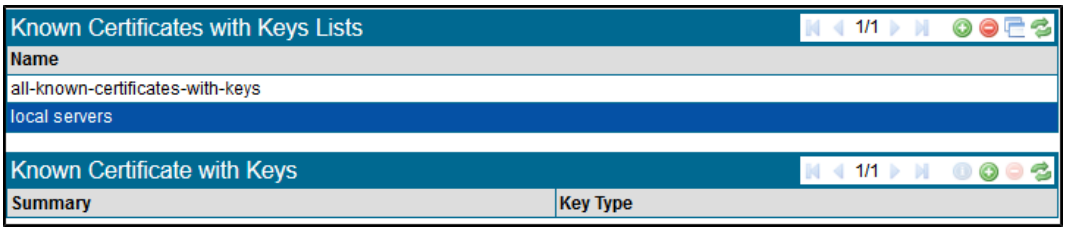

The figure shows the **PKI > Known Certificates with Keys List** window after a list called "local servers" has been added and saved. Initially this custom list has no entries as can be seen by the fact there are no entries in the **Known Certificates with Keys** area. To add entries to the list highlight the local-servers list and then click **Add** in the **Known Certificate with Keys** section.

To add keys and or certificates to the custom list, copy them from the **all-known-certificates-with-keys** list. The top section of the panel lists all the keys/certificates that are present in the all-known-certificates-with-keys list. Clicking on an item will highlight it, and clicking **Add** to **Custom List** will copy the item into the customer list.

In this example, the key/certificate for bluecoat.com has already been copied across. Once all the keys/certificates that need to be included in the custom list have been copied, click **OK**. At the bottom of the screen is a **Policy Changes** notification block with options to **Apply** or **Cancel** the change. Click **Apply** to complete the process and to save the CA to disk.

The ruleset for this example includes five rules.

Rulesets

Name	Rule Count
active-inline ruleset	4
ruleset	1

Ruleset Options

Default RSA Internal Certificate Authority:	(Not Set)
Default EC 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
Host Categorization IP Exclude List:	(Not Set)
HSM Failure Action:	Cut Through

Rules 1/1

Match Fields	Action	Comment
subject-domain[sslng unsupported sites]	Cut Through	Don't inspect as will break application
known-certificates-with-keys[all-known-certificates-with-keys]	Decrypt (Certificate and Key known)	local severs that we have key/cert for
subject-domain-list[webmail destinations]	Decrypt (Resign Certificate)	webmail systems
known-certificates[all-trusted-certificates]	Reject	Reject sessions to servers with expired certificates

The first rule uses the default **sslng-unsupported-sites** list to cut through traffic to any destinations that are in this list. Trying to inspect traffic to these sites will cause the application to break so the cut through rule is needed to prevent this.

The second rule uses the local-servers list to inspect traffic using known server key/certificate mechanisms. The third rule uses the **`webmail-destinations`** list to inspect traffic to webmail systems using certificate resign.

The fourth rule causes any SSL sessions to servers that have an expired server certificate to be rejected. The fifth rule is a "catch all" rule that means any SSL traffic that has not matched one of the preceding rules will be inspected using certificate resign.

Note Position of rules in the table matters as the list is processed from top to bottom. As shown the rule relating to expired certificates will not apply to servers in the local-servers list as this will be processed first. The up and down arrows can be used to alter the position of a rule in the **Rules** panel.

The final part of the process is to 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 to see the Segments information. To create a new segment click **Add** in the **Segments** table. The figure shows the segment configuration after it has been saved and activated. In this example you can see:

- The configuration allows the connection of an active security appliance, such as an IPS
- The configuration is a Fail To Appliance mode, so in the event of failure of the SSL Appliance, traffic will still flow through the active security appliance
- The session log is enabled for this segment

- The configuration allows the connection of one passive security appliance which receives a copy of the traffic being sent to the active appliance
- Generated flows containing decrypted traffic are marked by changing the src MAC address to the value indicated.

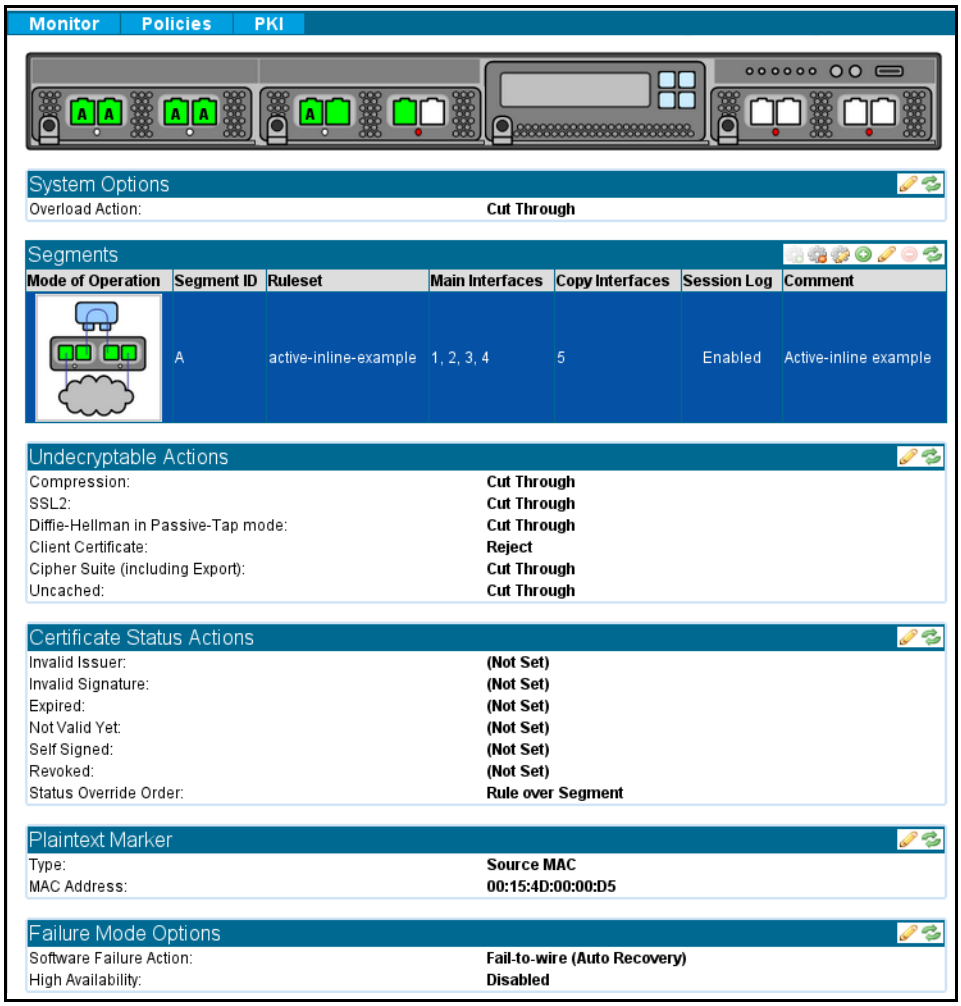

Work with a SafeNet Java HSM

A Hardware Security Module (HSM) provides additional security for storing cryptographic keys and certificates. The SSL Appliance is able to use a network-attached HSM appliance to store resigning CA keys, and to perform digital signature operations.

The SSL Appliance interacts with an HSM on its management interface. It exchanges signing requests and responses with the attached HSM appliance, over HTTPS. When mutually authenticated during the SSL handshake, the SSL Appliance sends resigning CAs data to the HSM; the HSM signs the data and returns the signature to the SSL Appliance.

An SSL Appliance can work with multiple HSM appliances, and multiple SSLV appliances can work with the same HSM.

In the event that policy a rule using an HSM to sign cannot work due to lack of response from the HSM, the attempt is logged, and the applicable policy action configured for HSM failure (for example, cut through, drop, or reject) occurs.

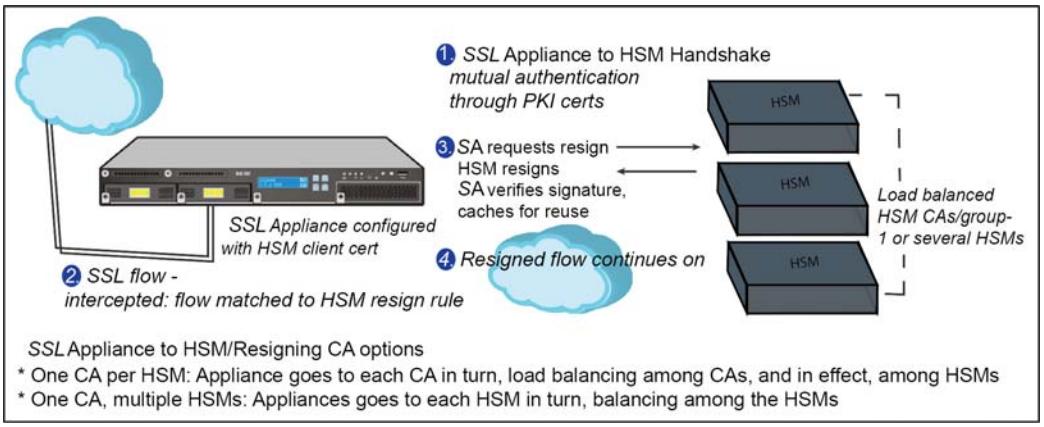

Cisco provides a Cisco HSM Agent and CLI to install on the SafeNet Java SP, which will be used to interact with Cisco appliances.

An SSL Appliance/HSM configuration has these basic steps, assuming the HSM is properly configured:

Step 1 Configure the client certificate(s) used to authenticate with the HSM (or HSMs).

Step 2 Configure an External CA(s) list used to authenticate the HSM.

Step 3 Configure the HSM, using the client certificate(s) and External CA List(s) configured above.

Step 4 Configure HSM resigning CA(s) using the HSM appliance configured above.

Step 5 (Optional) Add the HSM resigning CAs) configured above to HSM resigning CA load balancing group(s).

Step 6 Configure resign rule(s) in SSL Appliance policy using the HSM resigning CA load balancing group(s) configured previously, or the default **All HSM Resigning Certificate Authorities** group. If the default group is used, step 5 is not required.

Note Failure to apply changes might cause a manual HSM test to fail. For example, if you add an External CA to the External Certificate Authorities Lists but don't click **Apply** to complete the update, the test might fail due to using the incorrect External CA.

Adding an HSM

Before You Begin: PKI Basics

HSM validation requires an external certificate authority list on the SSL Appliance. The SSL Appliance will validate non-self-signed HSM certificates using the external CA list selected for that HSM in the **PKI > HSM Appliances** window. The appliance will validate the server certificate chain.

PKI Notes

- The CN or SAN in the HSM certificate must match the HSM appliance hostname or IP address for validation to succeed.
- The HSM must be configured to use 2048-bit or 3072-bit web server keys.
- TLS connections to the HSM are allowed only when using the following SSL cipher suites:
  - AES128-SHA256
  - AES256-SHA256
  - DHE-RSA-AES128-SHA256
  - DHE-RSA-AES256-SHA256
  - AES128-SHA
  - AES256-SHA
  - DHE-RSA-AES128-SHA
  - DHE-RSA-AES256-SHA

Add a Trusted Certificate

The SSL Appliance requires a trusted certificate for the communication channel with the HSM. The SSL Appliance will only validate a self-signed HSM certificate if it is a Trusted Certificate.

The communication channel certificate is the HSM server certificate which the SSL Appliance will use to verify and trust the HSM.

Highlight the **all-trusted-certificates** item in the **Trusted Certificates Lists** panel, then click the **Add** (plus sign tool) icon in the **Trusted Certificates** panel. The **Add Trusted Certificates** window displays.

Upload the .pem or DER certificate file, or paste its text into the appropriate tab, then click **Add**.

The new certificate will appear in the **Trusted Certificates** panel.

Note Make note of the identifying details such as the DN and the key alias.

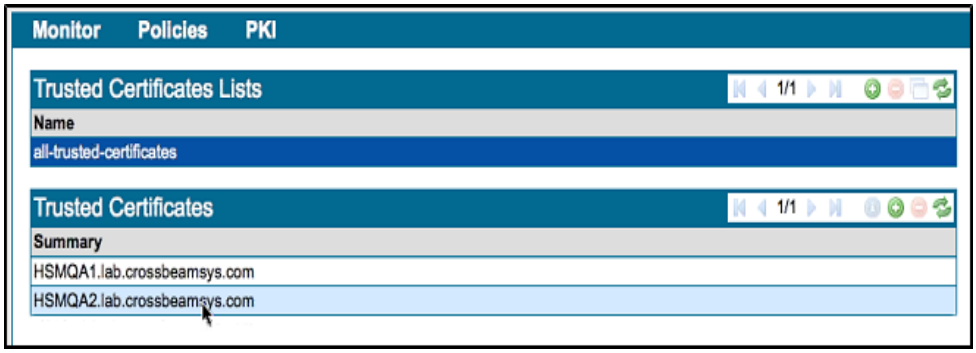

Add a Client Certificate

The SSL Appliance is a client to an HSM server, so the HSM must have a client certificate to authenticate the SSL Appliance. The SSL Appliance must have a client certificate for each HSM it interacts with. The client certificates must be available to create policy. The "Manage PKI" role is required.

PKI Notes

- Client RSA keys are restricted to 2048-bit and 3072-bit.
- Client Elliptic Curve keys are limited to keys generated on P-xxx, B-xxx, and K-xxx EC curves with primes of at least 244 bits.

An authorized user can

- Generate a client key and self-signed certificate:

Step 1 Generate a client key and self-signed certificate; click the Generate (rose) icon.

Step 2 Fill out the **Generate Certificate and Key** form.

Step 3 Click **Generate self-signed**.

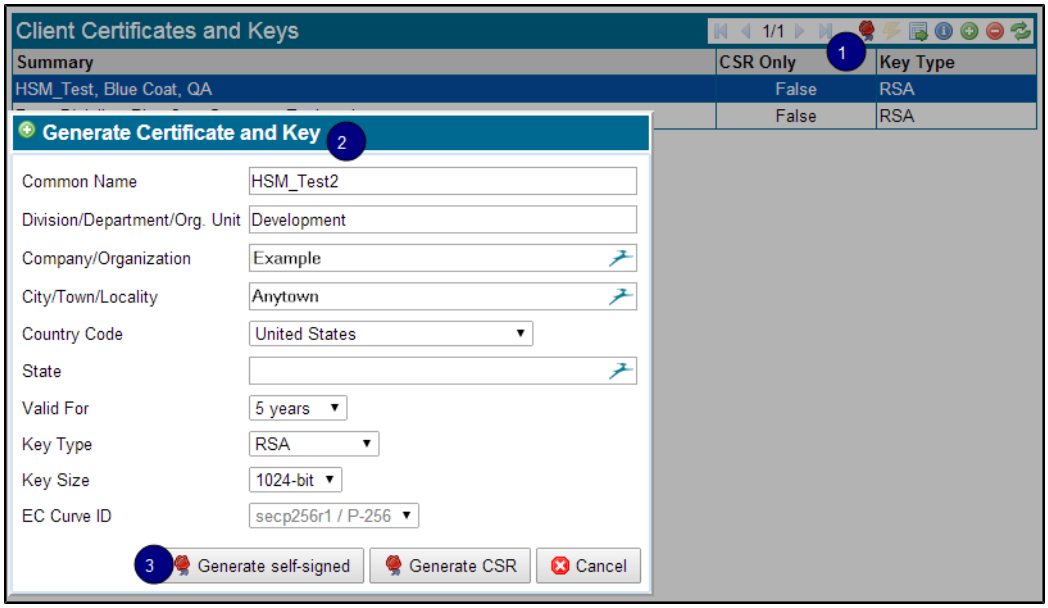

When the certificate has been created, you will see a confirmation message:

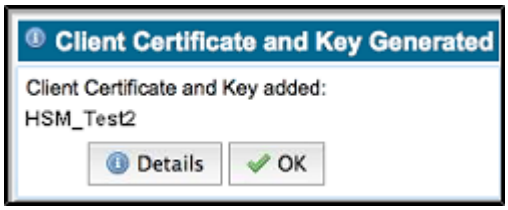

- Generate a client key and a CSR, and get it signed by a third party resigning CA.

Step 1 Click the Generate (rose) icon,

Step 2 Fill out the **Generate Certificate and Key** form.

Step 3 Click **Generate CSR**.

Step 4 Send the CSR to your third-party CA, and have it signed.

Step 5 Upload or copy in the signed certificate using the Install Certificate icon (lightning bolt).

- Import a client key and already signed certificate.

Step 1 Click Add (the plus icon).

Step 2 Upload or Paste in the already signed certificate.

Add an HSM

Use the **PKI > HSM Appliances** window to manage attached HSM appliance connections. A user must have a Manage PKI role to add, remove, and edit HSM appliances. Users with Manage PKI and Manage Policy roles can view configured HSM appliances.

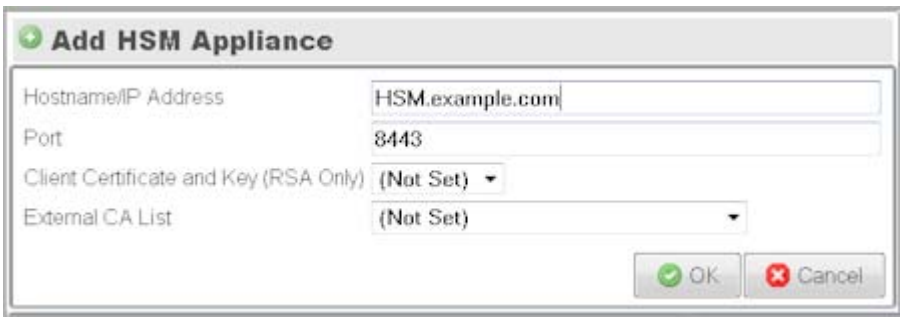

Hostname/IP Address: Enter the hostname that displays in the HSM-created certificates.

Port: Use the default 8443 port.

Client Certificate and Key (RSA Only): Select the client certificate created for the specific HSM.

External CA List: Select the External CA List to used to authenticate this HSM appliance.

Add Resigning Certificates

The SSL Appliance has a **HSM Certificate Authorities Groups** list of load-balancing groups on the **PKI > Resigning Certificate Authorities** window. The **all-hsm-certificate-authorities** list must contain all resigning certificates used by connected HSMs. The keys are stored remotely on the HSMs.

Creating, editing, and deleting HSM resigning CAs, and running self-tests, requires the Manage PKI authorization role.

A resigning certificate can be generated on the HSM with Cisco CLI. You can also create a CSR with the CLI, then sign the certificate off of the HSM.

Add HSM resigning CAs to the **PKI** store; they are initially added to the **all-hsm-certificate-authorities** list.

Add a new HSM Resigning Certificate Authority

Step 1 Highlight the **HSM Resigning Certificate Authorities Groups** you want to add an authority to, then click Add in the **HSM Resigning Certificate Authorities** panel. The **Add HSM Resigning Certificate Authority** window opens.

Step 2 You may upload or paste in the certificate on the appropriate tab.

- **HSM Appliance:** Select the HSM this resigning CA uses (created at **PKI > Add HSM**).
- **HSM Key Alias**: Enter the key alias configured for that resigning CA on the HSM appliance.
- **CRL URL**: Applies when a resigning CA CRL is published at a public URL.

Step 3 Click **Add**.

Step 4 Click **Apply** near the footer to save your changes.

Important Notes

- Immediately after adding a new HSM resigning CA to the PKI store, the new CA will appear in yellow in the list. Make sure to Apply the changes, and then use the Refresh tool in the header in order for the new CA to appear in green
- HSM Resigning CA Status Colors

The color of the resigning CA in the list gives you the status of that CA.

- Green color = HSM CA is in OK/active state
- Yellow color = HSM CA status has not been checked
- Red color = HSM CA is in failed state

Load Balancing Groups

The SSL Appliance uses round robin load balancing, distributing connections evenly across the array of HSM resigning CAs in the group, when the configuration includes a HSM resigning CA group, and an inspected SSL flow matches a policy HSM group rule. If each resigning CA uses a separate HSM appliance, the load balancing occurs over the HSM appliances. You can create a new subset group, if required. The SSL Appliance will load balance between the HSM resigning CAs in the group. An HSM CA can't be added to more than one user-defined group.

If an HSM in a group fails or is unreachable the SSL Appliance automatically adjusts the load balancing in the group, excluding the failed HSM.The SSL Appliance periodically checks the failed appliance, and when a check succeeds, the HSM is restored and returned to the load balancing group.

Create a new HSM authorities load balancing group

Step 1 On the **PKI > Resigning Certificate Authorities** window, select **Add** in the **HSM Resigning Certificate Authorities Groups** panel header.

Step 2 On the resultant **Add Resigning Certificate Authority List** window, provide a name for the list, then click OK.

Step 3 Highlight the new group in the **HSM Resigning Certificate Authorities Groups** panel

Step 4 Click Add in the **HSM Resigning Certificate Authorities** panel. The **Manage PKI Custom List Items** window opens.

Step 5 Click **Add to Custom List** and **Remove from Custom List** to create your list.

Step 6 Click **OK**.

Step 7 Click **Apply** near the footer to save your changes.

Run a Self-Test

Once you have the HSM connection configured, resigning certificates established, and a policy (ruleset) in place, you can manually run a self test to verify resigning. Click **Play** in the **HSM Resigning Certificate Authorities** header.

In the self-test, the appliance sends a signing request to the HSM, and then validates the returned signature. To pass the test, the HSM must have been contacted, the HSM resign operation succeeded, and the returned signature verified as valid.

Write HSM Configuration in Policy

The **Ruleset Options** panel (**Policies > Rulesets**) includes a **HSM Failure Option**. which defines what action the SSL Appliance will take when an HSM resign operation fails.

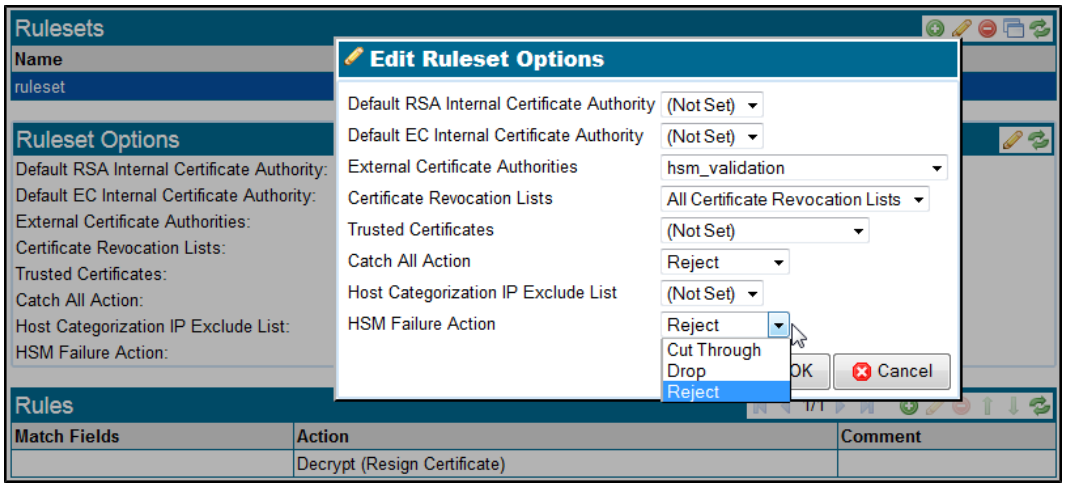

Step 1 Create and name a ruleset: click **Add** in the **Rulesets** panel header and name the ruleset.

Step 2 Click **Edit** (pencil tool) in the **Ruleset Options** panel header; the **Edit Ruleset Options** window displays, as shown in the figure. Here is the HSM selection to note:

- **HSM Failure Action**: Typically, you will Cut Through an SSL flow where the HSM resign operation has failed; you may also Reject or Drop the flows.

Step 3 Click **OK**.

Step 4 In the **Rules** panel, click **Add**. The **Insert Rule** window opens.

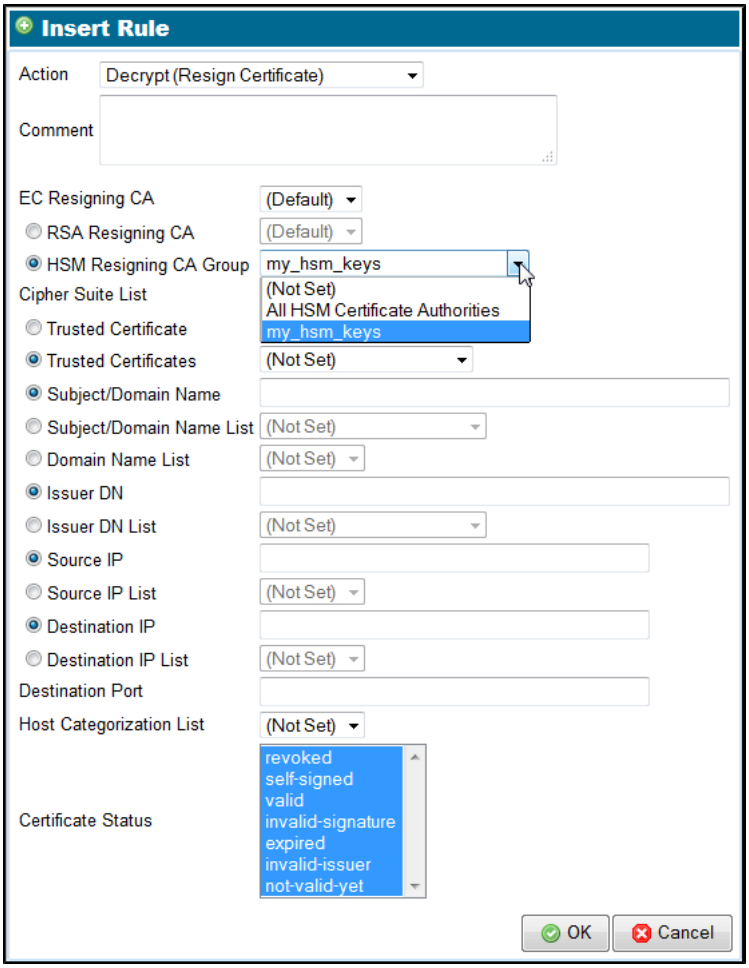

- **Action**: Select **Decrypt (Resign Certificate)**
- **EC Resigning CA:** Select the CA to resign flows signed with EC certificate authorities.
- **RSA Resigning CA:** Select the CA to resign flows signed with RSA certificate authorities.
- **HSM Resigning CA Group:** Select the **HSM Resigning CA Group**, then the HSM CA group which will resign HSM SSL traffic. Resigning traffic is load balanced across the CAs.

Step 5 Click **OK**. The window closes.

Step 6 On the **Rulesets** window, click **Apply** at **Policy Changes** in the footer.

When everything is set up, the SSL Appliance will inspect traffic which matches an inspection rule, resign it with the verified HSM resigning signature, and send the flow on to its destination.

HSM Logs

HSM interaction information is available in the logs (under the **Monitor** menu item).

System Log: View HSM failures as follows:

- HSM appliance failure
- HSM appliance recovery

SSL Session Log: View HSM signature failures and other errors as follows:

HSM network connectivity failure:	HSM or Agent can't be reached
HSM secure connection failure:	SSL connection failure between the SSL Appliance and the HSM
Invalid HSM response:	HSM response is corrupt, or can't be verified
Invalid HSM request:	HSM returned a HTTP 400 response
HSM operation internal error:	HSM returned an HTTP 404 or 500 response

HSM Diagnostics

The Policy section of a diagnostics report will contain the resigning CA fields and HSM Failure options in the rulesets, as well as the user-defined HSM resigning CA groups. The PKI section will include the HSM configuration, HSM resigning CAs, and the client certificates.

Note No diagnostic CLI is available for HSM connections for the SSL Appliance.

User Interface Overview

Introduction

This section provides details of all the facilities provided by the web-based user interface (WebUI) on the SSL Appliance device. 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 (Cisco recommends Internet Explorer and Chrome) and enter the hostname or IP address of the appliance in the address bar. The current IP address and hostname of the appliance can be viewed on the front panel LCD screen by pressing the bottom right button on the keypad until the Network option is displayed and then pressing the top left button. If the hostname has not been set yet, or if the hostname does not map to the IP address, the IP address must be used.

Configure the Browser

Accessing the web interface without the correct certificate installed in the web browser will cause the browser to display a warning message. This is the normal and correct behavior for the web browser. To prevent the warning message being displayed the browser needs to be configured to trust the certificate being used by the web server in the SSL Appliance.

There are two ways that the browser can be made to trust the SSL Appliance certificate. An SSL server certificate that is issued by a trusted CA can be loaded into the SSL Appliance this will be used by the internal web server and as it is issued by a CA that the browser trusts the browser will no longer generate a warning message.

The other method is to 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, page 6-72. If the browser generates warnings then you should consult your browser documentation for instructions on how to add the SSL Appliance certificate to the set of trusted certificates stored in the browser.

The next figure shows the warning produced by Chrome when accessing an SSL Appliance for the first time, followed by the Firefox warning.

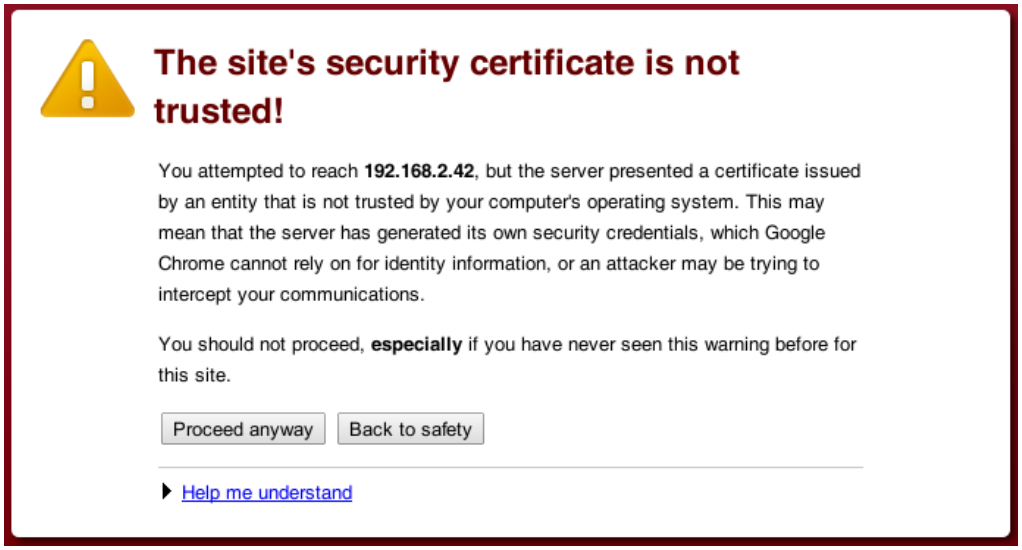

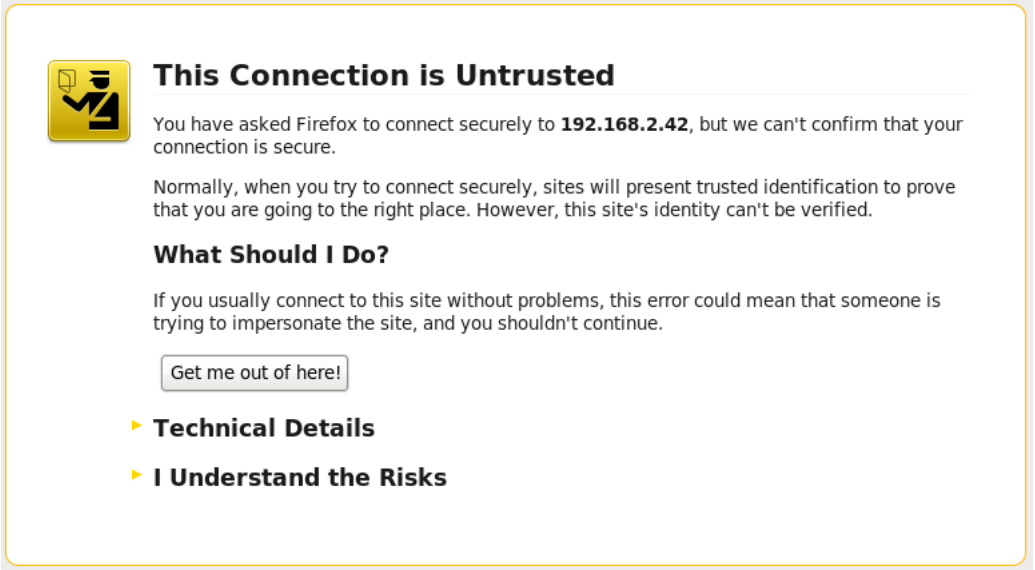

In both examples, the SSL Appliance had a management IP address of 192.168.2.42. In the case of Chrome clicking “Proceed Anyway” allows the browser to connect to the SSL Appliance. In the case of Firefox, click “I understand the risks” to access to screens that allow the certificate from the to be added 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 a user name and password are created and can then be used to log on to the system once the bootstrap phase is complete. See the *Getting Started Guide* for your appliance for details of the bootstrap process. Additional user names and passwords can be created using

the WebUI. Multiple users can be logged on at the same time. The system will rate limit login attempts to ten attempts in one hour to prevent attacks. The system will also timeout a session and then prompt the user to reenter the password before allowing access again.

You may inspect the EULA and software attributions without logging in.

Use the Main Screen

The management interface screens are laid out such that particular types of information are displayed in specific areas on the screen, no matter which screen you are looking at. The basic organization of the management screens is described below.

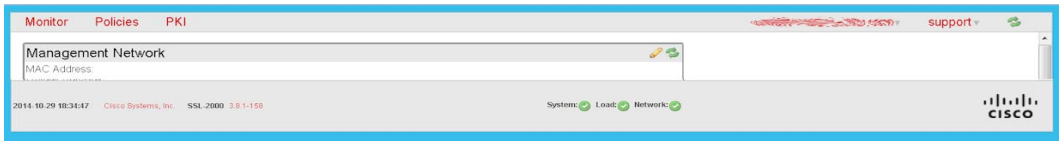

The top of the window contains five menus, a Refresh tool and, when a refresh is occurring, a spinner to indicate this. The menu items are explained in detail in later sections.

The bottom of the screen shows a status bar that is always present. It displays the following information:

- Current date in YYYY-MM-DD format
- Current time in HH:MM:SS format
- Copyright notice
- SSL Appliance Model Number
- Software version currently running on the system
- Icons showing current status for the System, Load, Network, and License.

The **System, Load, and Network** icons appearance varies as follows:


An error is present


A warning is present


Everything is fine

- The **License** icon appearance depends on the status of the license as follows:

	No valid SSL Appliance license is present, or the license has expired
	The installed SSL Appliance license expires within 30 days, and/or the Host Categorization license has expired
	A valid license, not expiring within 30 days, is installed

The active window or panel displays between the top and bottom bars, and is organized into panels; the example above shows the **Network Management** window. Each panel of the window has a title bar at the top and a set of tool icons at the right hand side.

The set of tools available varies by panel. Some the tools might be unavailable and grayed out, depending on how the panel is being used. Panels might also be empty, in which case only the title bar will be visible. An icon will be unaccessible (grayed out) if it doesn't apply to the selected item.

Overview of Common Tools

	Add; a new name. list, rule, CRL, and so on			Multipage tools; move to the first/last/next/previous screen
	Delete; a name, list rule, CRL, or so on			Refresh; reload the data
	Edit			Clone
	Information; view details			Acknowledge
	Move up			(action); such as download the Host Categorization database
	Move down			Disable (rule)

Note The multipage tool indicates which page from a number of pages of data the panel is currently displaying, along with tools for moving between pages within the panel, as explained above.
To move directly to a particular page, click on the numbers between the move forward and backward one tools, and then typing in the number of the required page.
Multipage panels have a built in multiplier that is used in conjunction with the number of rows value that is configured as the default (see Preferences, page 6-76). For example, 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.
Multipage panels are configured to display a maximum number of rows so the maximum number of pages that the panel displays is determined by the configured page size relative to the size of the log (see Preferences, page 6-76).

A display-only panel will have the **Refresh** tool, and might have the toggle **Auto Refresh** tool. The **Refresh** tool refreshes the data in the panel, while the toggle **Auto Refresh** tool turns on or off auto refresh.

Temperatures (Degrees ºC)

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
31	25	-43	-35	-48	-53	-60	54	37	41

Thermal margin sensors are reported as negative values which when increased to 0 will cause CPU to throttle down or halt

Some panels contain configuration data that can be edited; in this case there is an **Edit** tool in addition to the **Refresh** tool.

Date/Time

Date:	2014-2-4
Time:	01:43:13
Timezone:	Europe/Berlin
NTP Enabled:	False
Primary NTP:	
Secondary NTP:	

Panels might also be linked to other panels, so that an action taken in one panel will affect the related panel, as shown next.

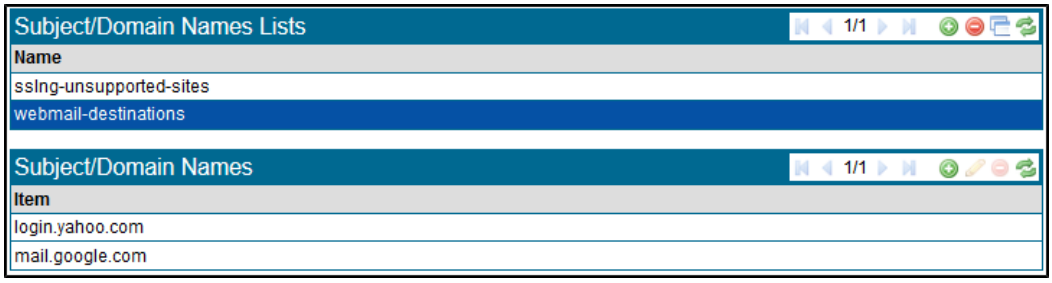

The top **Subject/Domain Names Lists** panel contains details of lists that are stored in the system and has tool icons allowing the following actions in addition to the **Refresh** action and multipage tools:

When a row in the top **Subject/Domain Names Lists** panel is selected the lower **Subject/Domain Names** panel will show the names contained in the list that has been selected and provides tools icons for you to manipulate them.

This covers the basic types of panel that are used by the system. Details on the specific panels used on different menus are covered in later sections of this document.

Monitor the System

The **Monitor** menu options provide details on the operation of the system and allow for the collection of diagnostic and debug information.

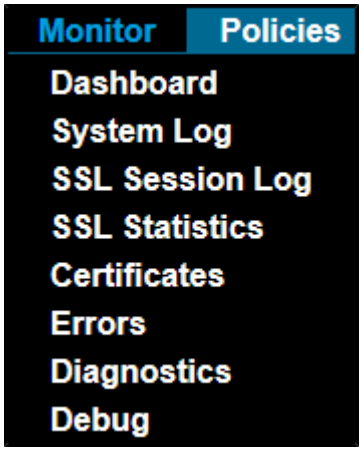

These options are described in detail below in the order in which they appear on the menu.

Monitor the Dashboard for Current Status

The dashboard display contains several panels containing different types of information, described below. The top of the dashboard display shows a graphical representation of the system that identifies which interfaces are being used by which segment, and indicates if the interface is active or not. Use the **Refresh** tool to present fresh data.

The image represents the physical configuration of the system so the number and types of Netmods (if applicable) matches the configuration of the system.

This figure shows the graphic for an SSL8200 system that has two 4 x 10/100/1000 copper Netmods installed. It shows that there is one active segments (A), and that one 10GigE port is active. All the ports that show green are up.

It shows that there is one active segments (A). Ports that show green are up.

Gather System Information

The **Segment Status** panel displays the status of currently active segments.

- The **Segment ID** is a unique identifier that enables this segment to be distinguished from other segments that might be present in the system.
- The **Main Interfaces** identify the physical ports that are being used by this segment. If any of the interfaces being used by the segment are currently down, the interface numbers will show in the **Interfaces Down** column.
- **Copy Interfaces** indicate active copy interfaces
- **Main Mode** indicates the operating mode of the segment.
- The **Failures** column will record any failure details.

Segment Status Tools

The tools specific to this window are:

- The **Manual Fail** tool forces a segment into a failed state. The state persists across a restart or reboot. **Manual Fail** is active if a segment is selected.
- The **Unfail** icon is active only if the segment is in a failure mode that requires manual intervention to clear the failure.

The background color for a segment row indicates if there are any problems with the segment.

Segments Status

Segment ID	Main Interfaces	Copy Interfaces	Interfaces Down	Main Mode	Failures
A	1, 2, 3, 4			Active-Inline-FTA	

The **Network Interfaces** panel has a row for every interface that is installed in the system.

- The maximum number of rows for an SSL2000 is 12 if it is fitted with three 4 x 1Gig Netmods.
- The maximum number of interfaces on an SSL8200 is 16.

•

The **Link State** column shows the speed that the link is operating at. For example, when a 1G interface is in use, it can operate at 10 Mbps, 100 Mbps or GigE rates.

Network Interfaces

Port	Type	Link State	RX Packets/Bytes	TX Packets/Bytes	RX Drops
1	1G	1G	18864475/15662763383	19085341/14955569986	0
2	1G	1G	18967611/14946890174	18755544/15616215820	0
3	1G	Down	0/0	0/0	0
4	1G	Down	0/0	0/0	0
5	1G	1G	19955564/14931573650	15165425/15282183930	0
6	1G	1G	15165425/15282183930	19955571/14931583506	0
7	1G	Down	0/0	0/0	0
8	1G	Down	0/0	0/0	0

Each row shows each port and its corresponding interface type, and the speed it is operating at, along with transmit and receive statistics.

Power Supply

Left Power Supply	Right Power Supply
On	Off

View the status for each power supply in the **Power Supply** panel:

- On: Power supply is on.
- Off: Power supply is off.
- Not installed: A power supply is not present in the system.
- Failed: The power supply has failed.
- Checking: The power supply status is unknown.

See Sensor Thresholds, page 8-1 for information about sensor threshold values.

CPU Load %

cpu	cpu0	cpu1	cpu2	cpu3	cpu4	cpu5	cpu6	cpu7	cpu8	cpu9	cpu10	cpu11	cpu12	cpu13	cpu14	cpu15
0.2	1.2	0	2.9	0	0	0	0	0	0	0	0	0	0	0	1	0

The current **CPU Load %** window shows CPU utilization as a percentage of the total capacity of the CPU. The number of CPUs present depends on the platform.

The utilization threshold ranges are: midrange: 60%, high range: 90%. Colors provide status information:

- Red: Critical
- Yellow: Caution

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
13276	11860	13876	11592	7237	13157	13157

The **Fan Speed** panel displays the current speed values in RPM for the fans in the system. Color provides additional information:

- Red: Fan failure
- Yellow: Lower caution RPM
- White: Power supply is not powered
- Not shown: Power supply is not installed

See Sensor Thresholds, page 8-1 for information about sensor threshold values.

Temperatures (Degrees ºC)

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
31	25	-43	-35	-48	-53	-60	54	37	41

Thermal margin sensors are reported as negative values which when increased to 0 will cause CPU to throttle down or halt

The **Temperatures** panel presents details of temperatures (in Celsius) and thermal margins for components within the system; this is an SV1800 example. Colors provide status information.

SSL2000 and SSL8200 Notes

- Red: Critical temperature exceeded
- Yellow: Caution: midrange temperature
- White: Power supply is not powered
- Not shown: Power supply is not installed

SSL1500 Notes

- Red: Critical temperature exceeded
- White: Power supply is not powered
- Not shown: Power supply is not installed

See Sensor Thresholds, page 8-1 for information about sensor threshold values.

Utilization %

Disk	Memory
51.0	6.0

View the percentage utilization of system memory and disk space on the **Utilization** panel. The utilization thresholds are: midrange: 60%, high range: 90%. Colors provide status information.

- Red: Critical
- Yellow: Caution

The **System Log** panel contains the most recently generated system log entries, this panel automatically refreshes. It is also available as a distinct window. See the next section.

View System Log Entries

Use the **System Log** screen to view all entries in the system log. The panel has the multipage navigation tools, as well as **Refresh** and **Search**. The **System Log** supports up to three GB of history.

System Log 1/606

Time	Process	Log
Jan 31 22:06:01	kernel	imklog 5.8.6, log source = /proc/kmsg started.
Jan 31 22:06:01	rsyslogd	[origin software="rsyslogd" swVersion="5.8.6" x-pid="667" x-info="http://www.rsyslog.com"] start
Jan 31 22:06:01	kernel	[0.000000] Initializing cgroup subsys cpuset

Colors provide status information:

- Red: Error
- Yellow: Warning

Data displayed includes license information (License, page 6-69 and System Status, page 4-6) as well as system processes.

Click the **Search** tool to bring up the **Filter on Process** pop-up, where you can filter log entries to display only entries created by a particular process. Valid inputs are the names of processes which appear in the process column in the panel.

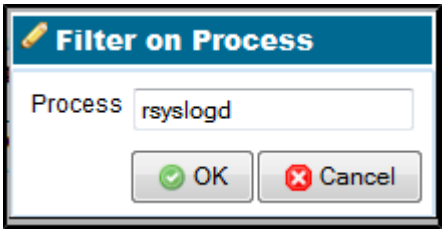

To cancel a filter, open up the **Filter on Process** window and delete the text in the input field, then click **OK**.

SSL Session Log

The **SSL Session Log** screen contains a single multipage panel enabling entries in the SSL Session log to be viewed. The **SSL Session Log** supports up to 32 million entries.

The log captures all flows which look like SSL, though some flows might not actually be SSL (a “false positive”). Each entry includes a flag which indicates if the session has been validated as SSL, and whether policy has been applied to the session. Use a post-processing tool to filter out these false positives.

The tools specific to this panel are:

	Export; brings up a window where you can specify the range of SSL session log entries to export.
	Filter on errors; causes the session log to only display entries for flows that were not inspected successfully
	No filter; causes the session log to revert to showing all entries.

SSL Session Log

Start Time	Segment ID	SrcIP:Port	DstIP:Port	Domain Name	Certificate Status	Cipher Suite	Action	Status
Mar 18 22:37:07.723	A	24.154.127.184:33387	23.210.249.115:443	sb.monetate.net	Valid	TLS_RSA_WITH_AES_256_CBC_SHA	Cut Through	Success
Mar 18 22:36:07.825	A	24.154.127.184:51898	74.125.28.104:443	Multiple domains	Valid	TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256	Cut Through	Success
Mar 18 22:29:25.054	A	24.154.127.184:33383	23.210.249.115:443	Multiple domains	Valid	TLS_RSA_WITH_AES_256_CBC_SHA	Cut Through	Success
Mar 18 22:29:18.565	A	24.154.127.184:33382	23.210.249.115:443	Multiple domains	Valid	TLS_RSA_WITH_AES_256_CBC_SHA	Cut Through	Success
Mar 18 22:28:49.863	A	24.154.127.184:33381	23.210.249.115:443	Multiple domains	Valid	TLS_RSA_WITH_AES_256_CBC_SHA	Cut Through	Success
Mar 18 22:28:36.421	A	24.154.127.184:51533	173.194.46.52:443	Multiple domains	Valid	TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256	Cut Through	Success
Mar 18 22:28:18.818	A	24.154.127.184:33379	23.210.249.115:443	Multiple domains	Valid	TLS_RSA_WITH_AES_256_CBC_SHA	Cut Through	Success
Mar 18 22:27:37.563	A	24.154.127.184:51891	74.125.28.104:443	Multiple domains	Valid	TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256	Cut Through	Success
Mar 18 22:25:07.776	A	24.154.127.184:52072	74.125.28.105:443	Multiple domains	Valid	TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256	Cut Through	Success
Mar 18 22:24:15.029	A	24.154.127.184:59475	74.125.28.106:443	Multiple domains	Valid	TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256	Cut Through	Success

Set the start and end date and time that the exported session logs should cover in the **Export** window. Click **Export** in the **SSL Session Log**, and the standard save file process on the browser will be invoked, which might automatically save the export file to a default location, or might prompt you to specify a location.

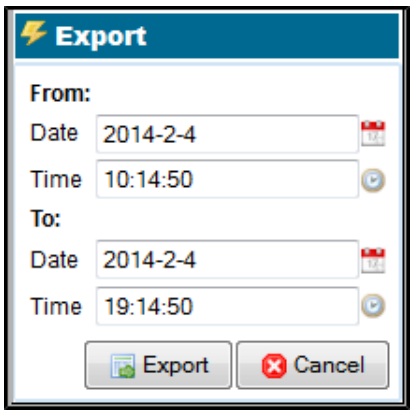

The saved file contains a set of .bin files and a file that contains the public certificates used in the SSL sessions captured in the session log. In order to view the session log data the .bin files must be processed with a tool to extract the data in a user readable form. The tool (SSL Sessions Tool) and tool documentation (sslsessions.pdf) are available on BlueTouchOnline (https://bto.bluecoat.com/) in Downloads. A Getting Started Guide is available on BTO Documentation.

The **Session Log** includes the following details for each SSL session that is 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
- Domain name of the SSL server accessed 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. So, the first row on page 1 is the most recent entry and the last row on the last page is the oldest

entry.The **View Details** tool is only active when a row in the **SSL Session Log** panel has been selected. Click it to open the **SSL Session Log Details** window and see s more details about the selected session. Clicking on the + or : symbol at the start of a line will expand or contract the level of detail displayed.

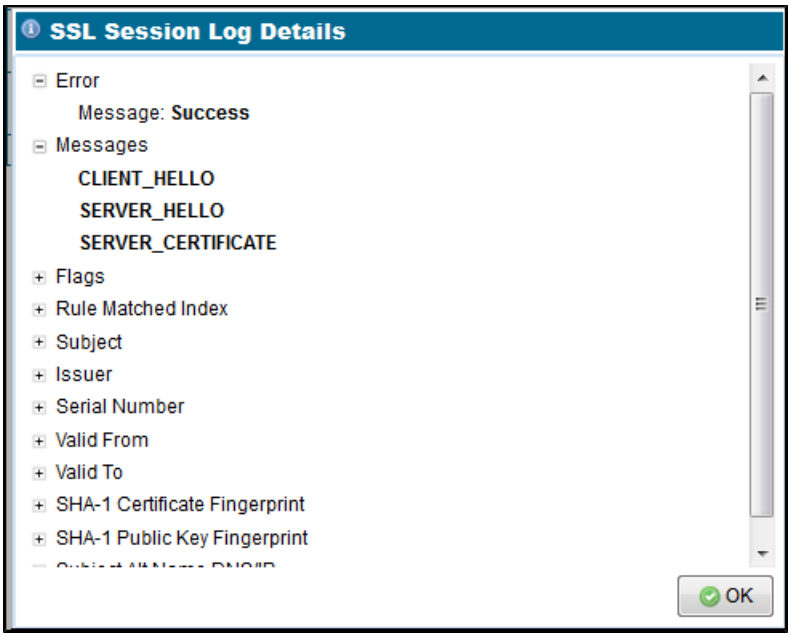

SSL Statistics

The **SSL Statistics** screen contains a single multipage panel for viewing the entries in the **SSL Statistics** log. The panel has the standard multipage navigation and **Refresh** tools. The **SSL Statistics** log supports up to seven days of history.

The next figure shows an example where page 1 of the available statistics information is being displayed. Statistics are collected every second and each row in the table holds the data for a collection interval. Apart from the **Detected** and **Decrypted** columns, all the counts are cumulative.

SSL Statistics 1/64

Timestamp	#Detected	#Done	#Ignored	#Decrypt	#Decrypt Done	#Error	Detected	Decrypt
Feb 12 14:19:05	13361	13301	809	11474	11418	18	60	56
Feb 12 14:19:04	13361	13300	809	11473	11417	18	61	56
Feb 12 14:19:03	13359	13299	809	11472	11416	18	60	56
Feb 12 14:19:02	13357	13297	809	11470	11414	18	60	56
Feb 12 14:19:01	13357	13297	809	11470	11414	18	60	56
Feb 12 14:19:00	13356	13296	808	11470	11414	18	60	56
Feb 12 14:18:59	13352	13294	808	11466	11412	18	58	54
Feb 12 14:18:58	13351	13293	808	11465	11411	18	58	54
Feb 12 14:18:57	13351	13277	808	11465	11395	18	74	70
Feb 12 14:18:56	13351	13277	808	11465	11395	18	74	70
Feb 12 14:18:55	13351	13277	808	11465	11395	18	74	70
Feb 12 14:18:54	13351	13277	808	11465	11395	18	74	70
Feb 12 14:18:53	13350	13276	808	11464	11394	18	74	70
Feb 12 14:18:52	13349	13275	808	11463	11393	18	74	70
Feb 12 14:18:51	13349	13275	808	11463	11393	18	74	70
Feb 12 14:18:50	13349	13275	808	11463	11393	18	74	70

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 might have been in that category over the one second period. Entries in the **Statistics** panel are ordered from most recent to oldest. So, the first row on page 1 is the most recent entry and the last row on the last page is the oldest entry.

Certificates

The **Certificates** window contains tabs for accessing details on invalid certificates that have been received by the SSL Appliance.The tabs show details for different types of invalid certificate states. See the next table for the tools specific to this window.

	Acknowledge
	Export; export details of all invalid certificates to a .csv file.
	Disable dumping invalid certificates into the system log.
	Enable dumping invalid certificates into the system log.

The figure shows the panel displaying details of all certificates that the system has seen which had problems of some description.

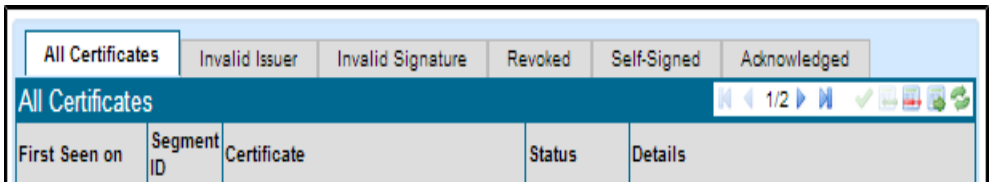

Click on the relevant tab to see details for specific types of invalid certificates; for example, the details of self-signed certificates that have been seen by the system.

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 noted. Once a certificate has been acknowledged it will appear on the acknowledged tab only. To acknowledge a certificate, select the certificate and then click on the tool. Acknowledged certificates will not be included in details on invalid certificates that are collected in the system log files.

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

Errors

The **SSL Error Counts** screen contains a single panel that shows SSL error counts for each active segment. Error counts are cleared when changes are made to the current ruleset, and policy is reset. The panel has the standard multipage controls in addition to **Refresh** and **Export**. Click **Export** to export the details of all errors to a .csv file.

Note An appliance functioning perfectly might have a non-zero SSL Error Count. An error count doesn't necessarily mean something is wrong.

SSL Error Counts 1/1

Segment ID	Code	Message	Count
A	0x34000017	Invalid MAC	1
A	0xc800014e	TCP queue processing timeout	1
A	0xc800034e	TCP queue processing timeout	4
A	0xc800044e	TCP queue processing timeout	19

The example shows a panel with a single invalid MAC address error, and multiple flows which ended without a FIN/RST sequence. There might 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 segment are reset to zero.

Diagnostics

Use the **Diagnostics** screen to specify the types of information to include in the diagnostic file, and generate the file.

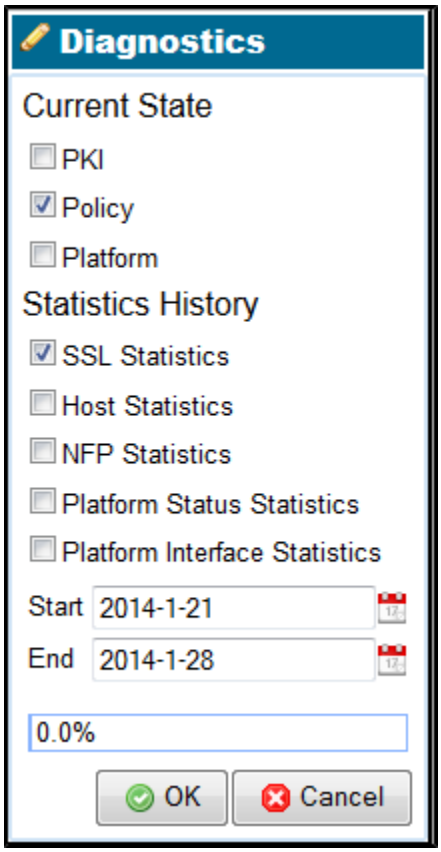

The example shows the window with **SSL Statistics** selected for inclusion in the diagnostic file. Checking the box next to an item will cause it to be included in the diagnostic file. The date fields can be used to limit the statistics/history data included in the diagnostic file. Click **OK** to create the file.

Caution Including the SSL Statistics and/or the Host Statistics, and/or the NFP statistics, might result in a large diagnostic file. Use these only if really required.

Debug

The **Debug** display presents NFE Network Statistics. The information on this screen is, as the name implies, primarily intended to assist with debugging issues with the SSL Appliance. Support personnel might ask for information from the debug screens when providing support. The **NFE Network Statistics** panels contain information that might be useful to a user in diagnosing configuration issues and some of the pages on the panel are described below. The **Debug** display supports up to seven days of history.

The **NFE Network Statistics** panel shows details of traffic to and from the Netronome Flow Engine (NFE) acceleration card(s) used in the appliance.

In an SSL2000 or SSL8200 the NFE card has two 10 Gbps links that connect to an Ethernet switch which in turn connects to the set of Netmods that provide the external interfaces on the SSL2000 and SSL8200.

For the SSL2000 there are two NFE links in the system; an SSL8200 has four NFE links, and will display two extra columns of data.

NFE Network Statistics 1/3		
	i/f 0	i/f 1
BadCRC	0	0
BadOctetsReceived	0	0
BroadcastFramesReceived	23775	23775
BroadcastFramesSent	23775	23775
CRCErrorsSent	0	0
Collisions	0	0
ControlFrameReceived	0	0
ExcessiveCollisions	0	0
FCReceived	0	0
FCSent	71437997	71437997
Fragments	0	0
Frames1024toMaxOctets	3136667	5244443
Frames128to255Octets	699907	701201
Frames256to511Octets	220317	226869
Frames512to1023Octets	227059	244350
Frames64Octets	0	0

NFE Network Statistics 2/3		
	i/f 0	i/f 1
Frames65to127Octets	4335636	3455977
GoodOctetsReceived	1894206543	4869814510
GoodOctetsSent	3400165410	3368384592
GoodUnicastFramesReceived	3598608	5361243
InRangeLengthErrorsReceived	0	0
Jabber	0	0
LateCollisions	0	0
MulticastFramesReceived	129060	129060
MulticastFramesSent	129060	129060
OutOfRangeLengthErrorsReceived	0	0
Oversize	0	0
ReceiveFIFOOverrun	0	0
RxErrorFrameReceived	0	0
SentDeferred	0	0
SentMultiple	0	0
SymbolErrorReceived	0	0

NFE Network Statistics 3/3

	i/f 0	i/f 1
TxBackoff	0	0
TxCarrierSenseErrors	0	0
TxExcessiveDefer	0	0
Undersize	0	0
UnicastFramesSent	4725363	4216526

The NFE card used in the SSL1500 has eight 1 Gbps links that connect to the external interfaces on the appliance via the Fail To Wire hardware and shows eight columns of data.

Configure Segments and Policies

Use the **Policies** menu to configure segments and define policies and rules that determine how SSL traffic is handled, and which SSL traffic is inspected. Overall, policies consist of rulesets, segments, and lists.

The top two **Policies** options are for configuring **Rulesets** and **Segments**, while the remaining options let you configure lists that can be used within **Rulesets.** These options are described in detail below in the order in which they appear on the menu.

Common Policies Tools

	Add list			Multipage tools
	Delete list			Refresh
	Edit			Clone

Configure Rulesets to Handle SSL Traffic

Rulesets contain the rules and policies that control how SSL traffic is handled. They are associated with one or more segments. Rulesets can also exist unassociated with any segment. Rules are followed in a top-to-bottom hierarchy, with the first rule matched determining the policy action.

The **Rulesets** display contains three panels. The lower two panels display information which depends on the row selected in the first panel.

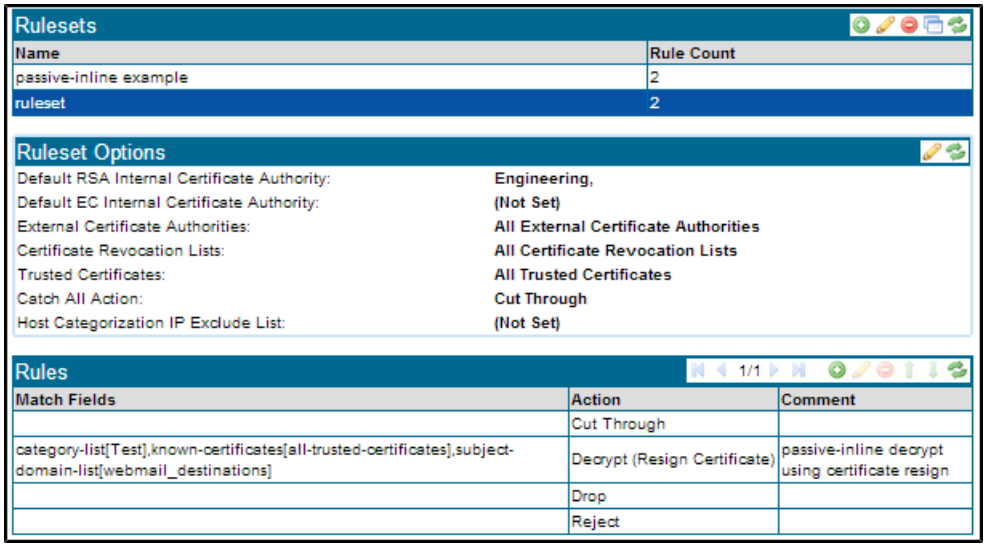

The figure shows the **Rulesets** panel with two existing rulesets. Each ruleset occupies one row in the table. The right hand column shows the number of rules that are currently within that ruleset. Tools let you **Add**, **Remove**, or **Clone** a ruleset. The **Remove** and **Clone** tools will be grayed out unless an entry in the table is selected.

If you select the **Clone** tool, the **Clone Ruleset** window displays, where you can configure the ruleset clone, as shown next.

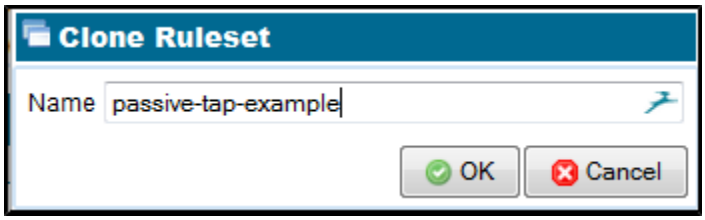

To view **Ruleset Options** and **Rules** information, select a ruleset entry in the **Rulesets** panel. To do so, click an entry; this will highlight the entry in the **Rulesets** panel. The **Ruleset Options** panel will expand and become active. The **Rules** panel displays the rules that exist within the selected ruleset.

Use the **Edit Ruleset Options** panel to configure the ruleset settings.

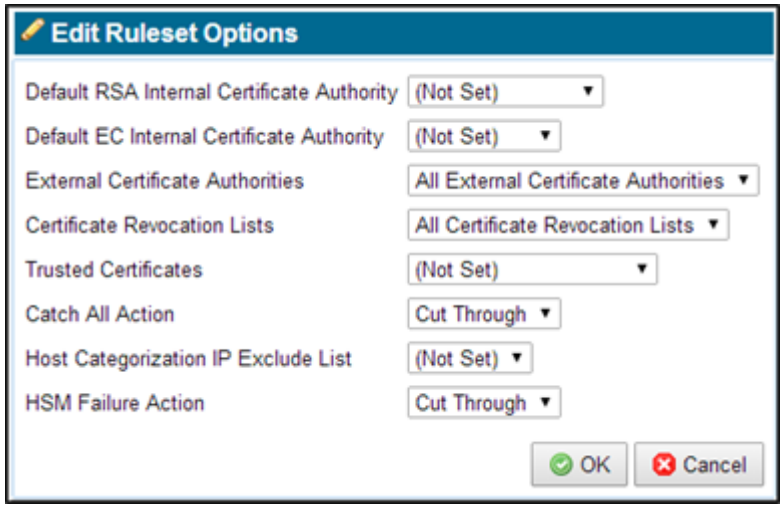

The **Edit Ruleset Options** are:

- **Default RSA Resigning Certificate Authority**: Used for "Decrypt (Resign Certificate)" rules where no RSA resigning CA is specified
- **Default EC Resigning Certificate Authority**: Used for "Decrypt (Resign Certificate)" rules where EC resigning CA is specified
- **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
- **Host Categorization IP Exclude List**: Selects the Host Categorization IP Exclude list as the list to check against when SSL sessions are processed by rules within this ruleset. See Host Categorization Lists, page 6-36.

- **HSM Failure Action**: Determines how to process an intercepted flow when the required HSM signature has failed.

The **Rules** panel, the bottom panel, displays the rules defined in the ruleset being edited. Click **Add** to display the **Insert Rule** window.

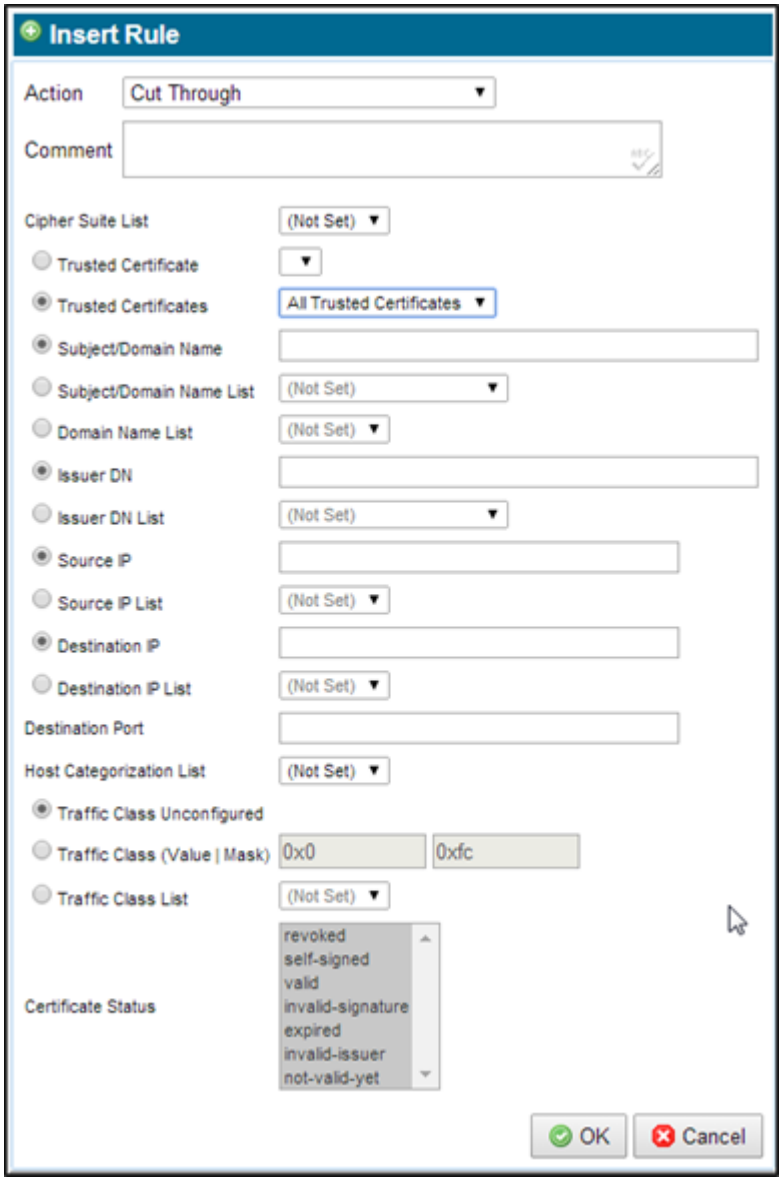

Use the **Action** drop down menu to select of the type of rule to create. When a selection is made, the window updates to display only fields relevant to the type of rule selected.

See Policy Rulesets, page 3-15 for an explanation of the parameters that can be configured for the different types of rules. For example, if **Cut Through** is selected, the **Insert Rule** window will appear as shown above.

Note 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 from the first rule in the table (top row on page 1) to the last rule in the table (bottom row on last page) so if a more generic rule occurs in front of a more specific rule, the generic rule will be encountered first and will always be used. An example will make this clear:

Rules 1/1

Match Fields	Action	Comment
subject-domain-list[sslng-unsupported-sites]	Cut Through	cut through sites cannot inspect to
subject-domain[webmail systems]	Decrypt (Resign Certificate)	webmail systems
known-certificates-with-keys[all-known-certificates-with-keys]	Decrypt (Certificate and Key known)	inspect local server traffic
known-certificates[all-trusted-certificates]	Reject	reject sessions to servers with expired certificates
	Decrypt (Resign Certificate)	inspect using certificate resign

The fourth rule is highlighted. It prevents any SSL sessions to destinations that have an expired SSL server certificate. The third rule causes traffic to destinations that are in the webmail list to be inspected. As the third rule will always be processed before the fourth rule traffic to any system in the webmail list will be inspected even if that system has an expired SSL server certificate.

In order to ensure that traffic is not allowed to a system in the webmail list if it has an expired server certificate the position of the highlighted rule needs to be changed so that it comes before the rule inspecting traffic to systems in the webmail list. To correct this, select the highlighted rule, then click the Up arrow to move it up in the table so that it is positioned above the rule inspecting traffic to systems in the webmail list.

Notes

- 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.
- The following server certificate match fields will be ignored when matching SSL flows using Anonymous Diffie-Hellman cipher suites:
  - **Subject/Domain Name**
  - **Subject/Domain Name List**
  - **Domain Name List**
  - **Issuer DN**
  - **Issuer DN List**
  - **Certificate Status**

Disable a Rule

You can disable a rule within a ruleset. When creating or editing a rule, the **Enabled** option is selected by default; the rule is active (and its location in the ruleset matters as usual). When the option is cleared, the rule is not processed.

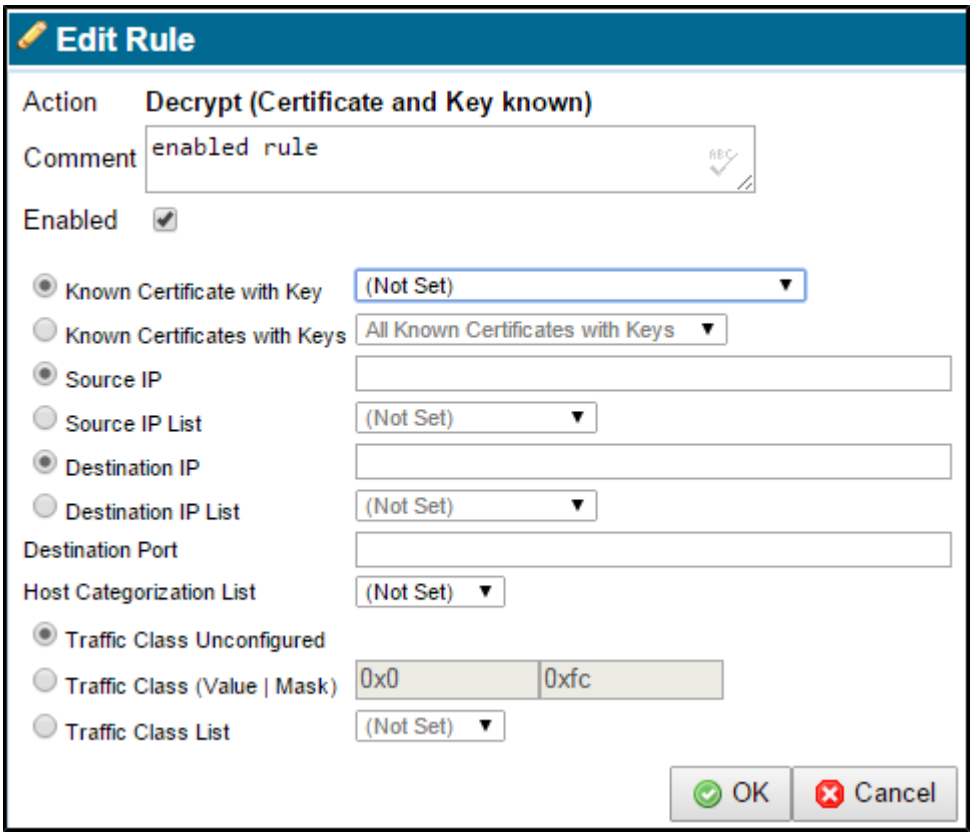

The setting is also shown per rule in the **Rulesets > Rules** panel, as True (enabled) or False (disabled) in the **Enabled** column. You can enable or disable a rule in the **Rules** panel. Select the rule and click the green check mark or red X icon at the top of the panel to enable or disable. Click **Apply** in **Policy Changes** at the bottom of the screen to save the change.

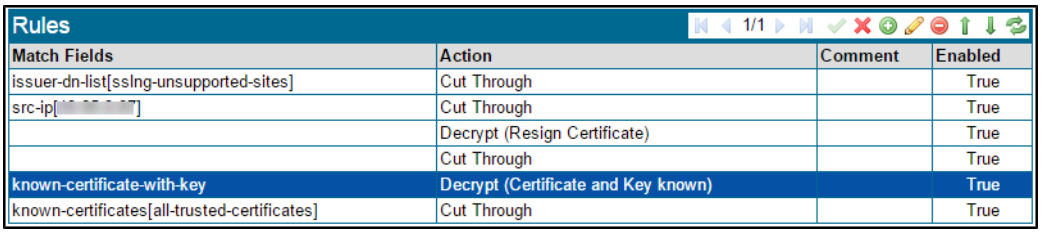

Rules 1/1

Match Fields	Action	Comment	Enabled
issuer-dn-list[sslng-unsupported-sites]	Cut Through		True
src-ip[ ]	Cut Through		True
	Decrypt (Resign Certificate)		True
	Cut Through		True
known-certificate-with-key	Decrypt (Certificate and Key known)		True
known-certificates[all-trusted-certificates]	Cut Through		True

In most situations, all rules should be set to True. If you are debugging a ruleset, you might use the False setting (that is, deselect **Enabled** for that rule), applying it to one rule at a time.

When a rule is disabled, its background display is yellow.

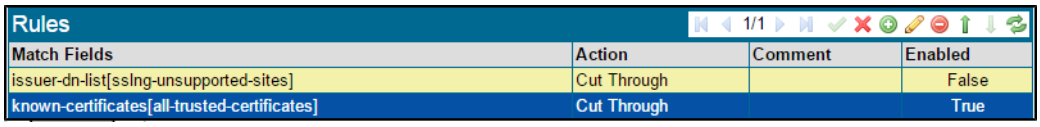

Rules 1/1

Match Fields	Action	Comment	Enabled
issuer-dn-list[sslng-unsupported-sites]	Cut Through		False
known-certificates[all-trusted-certificates]	Cut Through		True

Configure Receive Interfaces with Segments

A Segment is a grouping of interfaces which receives a network flow. The **Segments** display contains a graphical display of the system and six panels. The information displayed on the lower four panels depends on the row selected in the second panel

The first figure is an example of the graphic for an SSL2000 device. The graphic is dynamically created so it will reflect the set of interfaces that are installed in the box, in this case the unit has three 4 x 10/100/1000 Netmods installed.

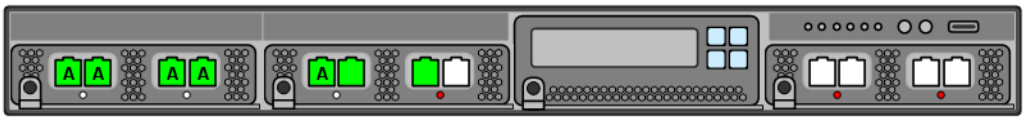

The next figure is an example of the graphic for an SSL1500device.

- Any interface that does not have a letter is not currently being used by an active segment.
- Any interface that shows green indicates that the relevant link is up.
- Deactivating an active segment releases the external interfaces used by that segment. those interfaces become available for use by other segments.

Activate a segment to use it. You will see a series of screens where you select interfaces for each required setting.

System Options

Use **System Options,** the first panel on the **Segments** screen, to configure the default action that the system should take if it is overloaded or if power is lost. In the example shown, the **Overload Action** is to cut through traffic; the other **Overload Action** options are to drop or reject the traffic.

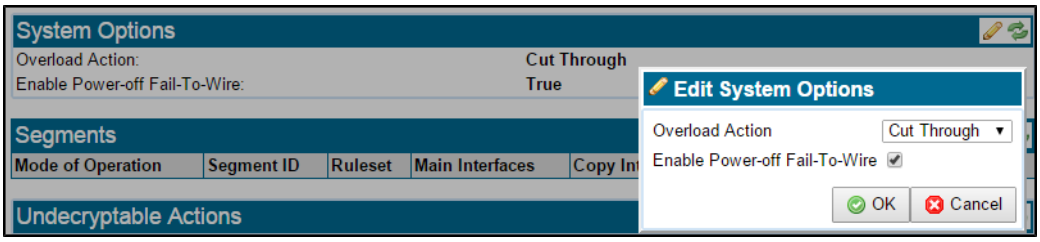

Enable Power-off Fail-To-Wire determines appliance behavior when the hardware is powered off.

- When **Enable Power-off Fail-To-Wire** is selected, on power-off, all traffic will be redirected from the incoming port to the paired port. This is the default setting.

- When **Enable Power-off Fail-To-Wire** is deselected, on power-off, traffic is redirected into the appliance instead of the paired port. All traffic is dropped.

Note During appliance startup, the interfaces might become active before the SSL inspection policy is applied. This lets traffic pass through the appliance without begin inspected. Use of a USB key or a keypad pin might allow this condition to exist for an indefinite period of time.

Append Resigning CA Chains to Resigned Certificates enables including the resigning CA certificate chain in resigned SSL sessions. This allows SSL clients to validate resigned certificates without auto-downloading the resigning CA certificate chain. Here is the basic procedure:

Step 1 On the **Segment > System Options** panel, check the **Append Resigning CA Chains to Resigned Certificates** option. The SSL Appliance will include the resigning CA certificate chain (configured in the PKI store) in the SSL session.

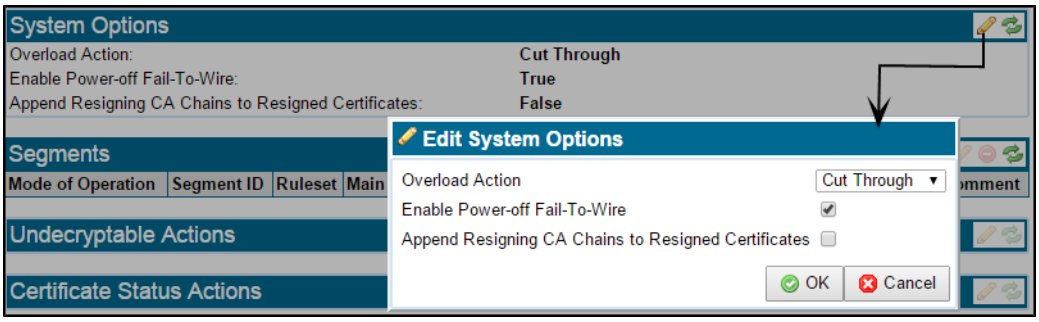

Step 2 On the **PKI > External Certificate Authorities** window, add all CAs from the resigning certificate chain to the **External Certificate Authorities** list. Once certificates have been added to the default **External Certificate Authorities** list, optionally create a new **External Certificate Authorities List**, and add the intermediate CAs which are included in the chain.

Step 3 On the **PKI > Resigning Certificate Authorities** window, add or edit a resigning certificate, **Local** or **HSM**. Select the required **Certificate Chain External CAs**.
Local CA Example:

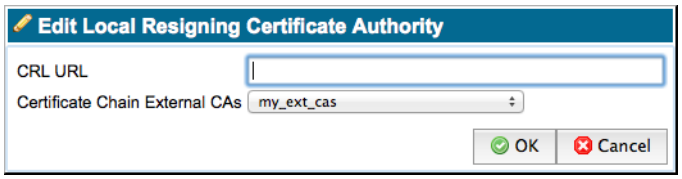

HSM CA Example:

Click **OK** (on an **Edit** window) or **Add** (on an **Add** window), then **Apply** the changes.

Step 4 Verify the CA chain. On the **PKI > Resigning Certificate Authorities** window, highlight the resigning CA, then click the **Test Certificate Chain** icon (chain link).
If the CA chain is complete, you will see a "Complete certificate chain is present" message.
If the CA chain in incomplete, you will see a "Incomplete certificate chain, first missing CA: <name>" message. Add the missing CA to the **External Certificate Authorities** list.

Step 5 Configure a new segment with a ruleset using the appended resigning CA.

Notes:

- During policy activation, the appliance will load the certificate chain for each active resigning CA from the External CAs.
- If a full certificate chain is not found for a resigning CA, a message will appear in the System Log, which identifies the first missing CA. The SSL appliance will load the partial CA chain and include it with resigned certificates in inspected SSL sessions.

Segments

The **Segments** panel (second from top) contains a row for each segment that is configured in the system. In addition to the **Add**, **Edit**, **Delete** and **Refresh** tools, it includes **Activate**, **Deactivate**, and **Edit Copy Mode** tools. A segment combined with a ruleset creates policy.

See Deployment Modes, page 3-7 for details of the modes of operation that can be selected for a segment when it is created. See Segment Policies, page 3-14. Passive-Tap Mode, page 3-8, Passive-Inline Mode, page 3-10, and Active-Inline Mode, page 3-11 provide examples of how to configure segments using the **Segments** panel. Here is an overview of the process:

View Segment Information

Once a segment definition exists in the **Segments** panel it can be selected by clicking on it. Once selected, the lower panels on the screen display information relevant to the selected segment.

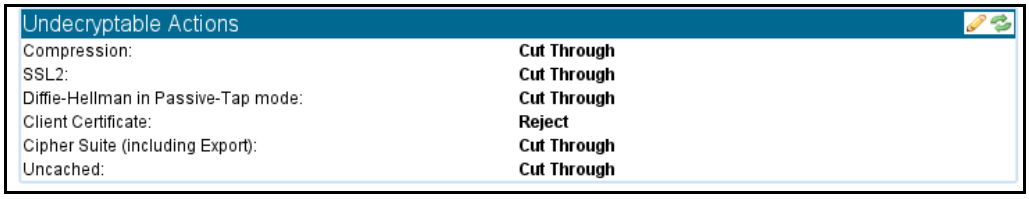

Use the **Undecryptable Actions** panel to control how SSL sessions on this segment that cannot be decrypted are handled. The panel has **Edit** and **Refresh** tools. Click the **Edit** tool to open the **Edit Undecryptable Actions** window and select the action to be take when a session is not decryptable for the specific reason. An SSL session cannot be decrypted for 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**: 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 in-line).
- **Client Certificate**: The use of client certificates in some situations can prevent an SSL Session being inspected. This action is applied when such a session is present.
- **Cipher Suite (including Export)**: The system does not support all possible SSL cipher suites: this action is applied when a cipher suite that is not supported is used by an SSL session.
- **Uncached**: An SSL session established using session re-use can only be inspected 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.

Certificate Status Actions

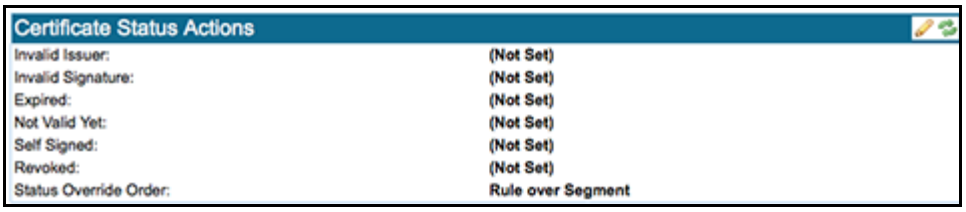

The **Certificate Status Actions** panel which lets you control how the system deals with SSL sessions on this segment that have particular states in the server certificate used for the session. The possible actions are: **Not Set, Cut Through, Drop,** and **Reject. Not Set** means that the particular status will be ignored.

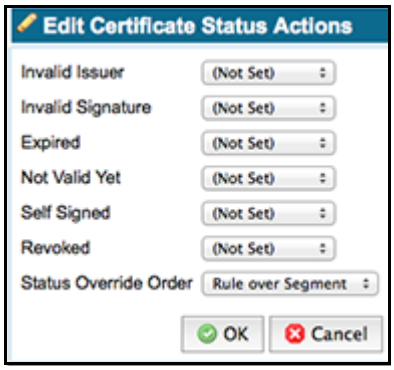

Use the **Edit Certificate Status Actions** window to configure the **Status Override Order**. This option determines whether or not the segment settings in this box take precedence over any settings in rules within the ruleset used by this segment. The options are **Rule over Segment** and **Segment over Rule**.

Plaintext Marker and Failure Mode Options

Use the **Plaintext Marker** panel and the **Failure Mode Options** panels to configure the failure mode and High Availability (HA) options. **Edit** and **Refresh** tools are available.

Click **Edit** in the **Plaintext Marker** panel to open a window where you control how generated TCP flows containing inspected traffic are marked.

There are two reasons for marking these flows:

- An attached passive security appliance might wish to be able to determine which traffic that it receives has been decrypted by the SSL Appliance and which has not. Configuring marking means the SSL Appliance will mark all generated flows and the attached appliance can use the marker to distinguish between inspected and non inspected traffic.
- If the SSL Appliance is configured to operate in Active-Inline mode then marking MUST be enabled as the SSL Appliance needs to be able to distinguish between inspected and non inspected traffic when it 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

You can also edit the **Appliance Feedback Options** to modify the **Feedback Timeout**. The **Feedback Timeout** determines how long the SSL Appliance waits for a response before canceling a request and interrupting the SSL flow. Selecting the **Extended** timeout allows a more time-consuming request, such as one to the cloud, to complete. The Default is 1 second. The Extended period is 5 seconds.

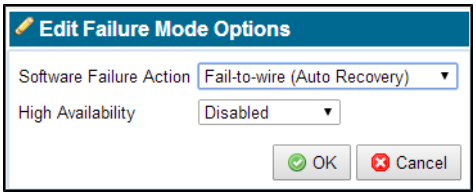

Clicking **Edit** in the **Failure Mode Options** panel produces the **Edit Failure Mode Options**, where you can configure how the system deals with software failures. The options, listed below, determine how this segment will behave in the event of software failure:

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

The High Availability mode options 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.

Translate VLAN IDs with VLAN Mappings

Use **VLAN Mappings,** the final panel on the **Segments** screen, to translate VLAN tags between ports. VLAN translation ensures traffic visibility as a packet traverses the network; it can be used as a way to distinguish each hop. When a segment using VLAN is selected in the **Segments** panel, the configured VLAN information displays in the **VLAN Mappings** panel.

About VLAN Configurations

Virtual LANs (VLANs) are *logical* network segments that allow hosts to communicate, regardless of physical network location. The benefit to this is that clients can be separated logically—based on organizational unit, for example— rather than based on physical connectivity to interfaces. The SSL Appliance treats VLAN interfaces identically to traditional physical LAN interfaces. VLAN segments are defined on the switch. The network administrator specifies which ports belong to which VLANs.

Traffic is processed as usual unless **VLAN Translation** is enabled. VLAN mapping does not separate tagged traffic, as a switch does.

The following diagram illustrates a simple VLAN configuration. The **VLAN Mappings** panel shows the configuration. The example shows an Active-Inline segment. The two bump-in-the-wire ports are NET-1 and NET-2. The two ports feeding traffic to the active tool are APP-1 and APP-2. All traffic flows in on NET-1, out to the security tool on APP-1, back in from the security tool on APP-2, and then out to the network on NET-2 (or vice versa for traffic going the other direction). Logical ports are assigned during the segment activation.

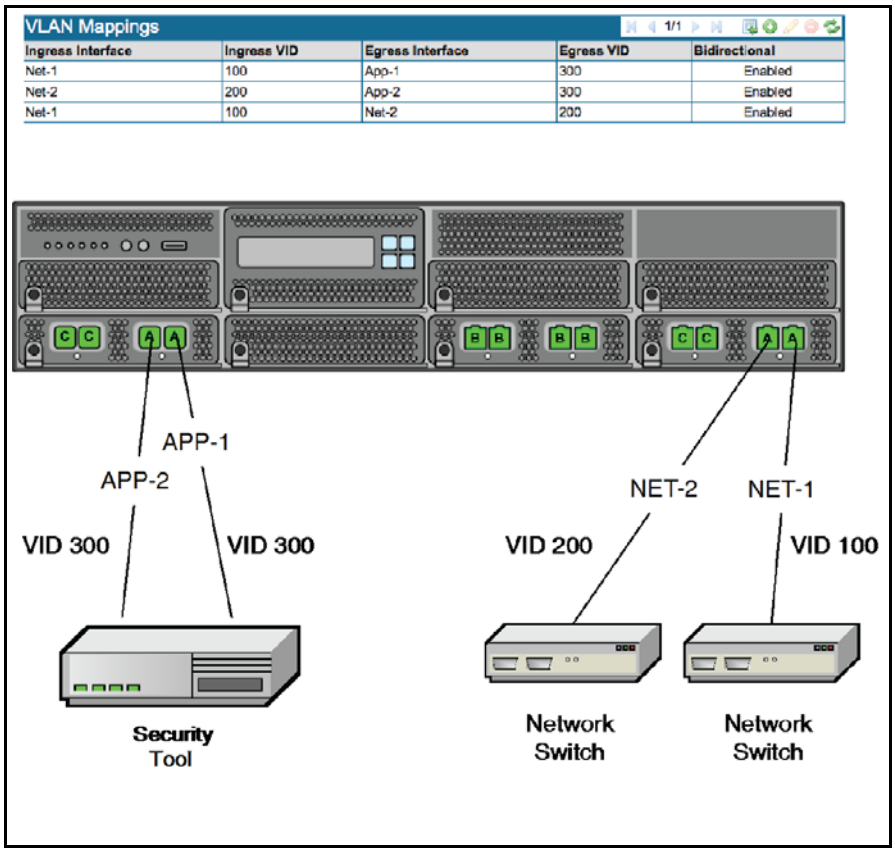

VLAN Notes

- Untagged packets remain untagged.
- Unmapped VLANs pass through unmodified. For example, unmapped packets which pass through the same interface as mapped packets which are being translated, remain unchanged.
- Packets transiting the same segment multiple times on different VLANs, whether or not VLAN translation is in effect, might be dropped or result in an SSL flow being processed incorrectly.
- Having the same flow occur on multiple segments with or without VLAN IDs will not cause a problem, as long as the flow occurs only once on each segment.
- VLAN mapping does not support stacked VLANs.
- The number of Net-x and App-x selections you must make depends on the mode of operation. For example, in an Active-Inline fail to appliance configuration, selecting Net-1 determines the remaining interfaces (App-1 will be the port paired with Net-1, and the remaining pair will be App-2 and Net-2).
  The most complex case is an Active-Inline Asymmetric fail to network configuration, where you must select the Net-1, Net-1 2, App-1 and App-1 2 translation interfaces.
- In a Passive-Inline deployment, the appliance (App-1) port is handled like a copy port, so a VLAN tag can be mapped between the two network ports, but there is no additional translation within the appliance. For example, if the SSL Appliance in Passive-Inline mode is translating between VLAN 10 on Net-1 and VLAN 20 on Net-2, traffic seen by the appliance will have a VID of 10 for traffic going out on Net-1 and a VID of 20 from traffic going out on Net-2.

- Translated VLAN IDs are not shown in packet captures which are taken on the appliance.

Example Basic Segment Configuration with VLAN Translation

This is a basic Active-Inline configuration, which includes VLAN translation. It provides an overview of configuring a segment, but the details will vary with the deployment. Make sure to know precisely which network connections carry traffic with VLAN tags when using VLAN translation.

Step 1 Add a segment and choose the operating mode.

In the **Segments** panel header, select **Add**. The **Add Segment** window displays.

- To select the **Mode of Operation**, click **Edit**. The **Select Mode of Operation** window displays.
- Select the mode of operation, either clicking on a graphic, or making selections from the **Main Mode** and **Asymmetric Sub-mode** menus.

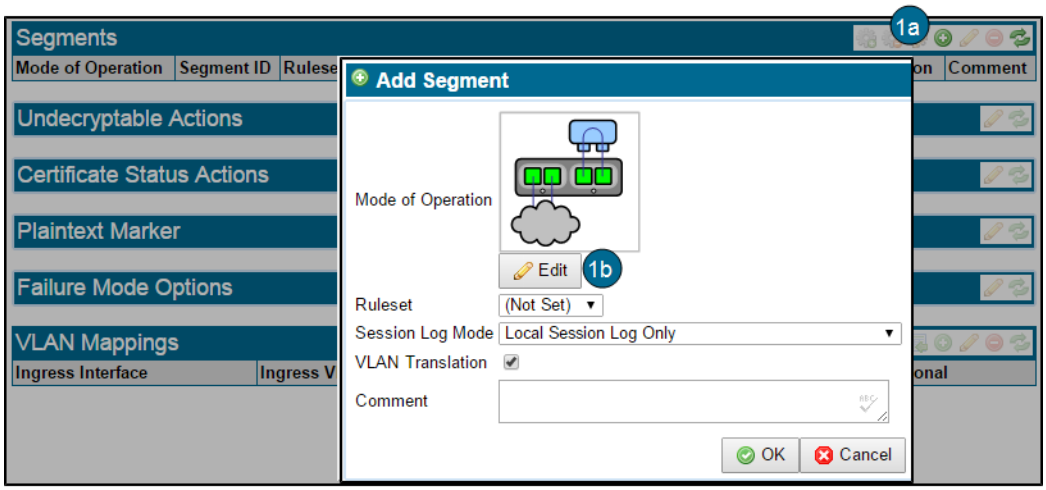

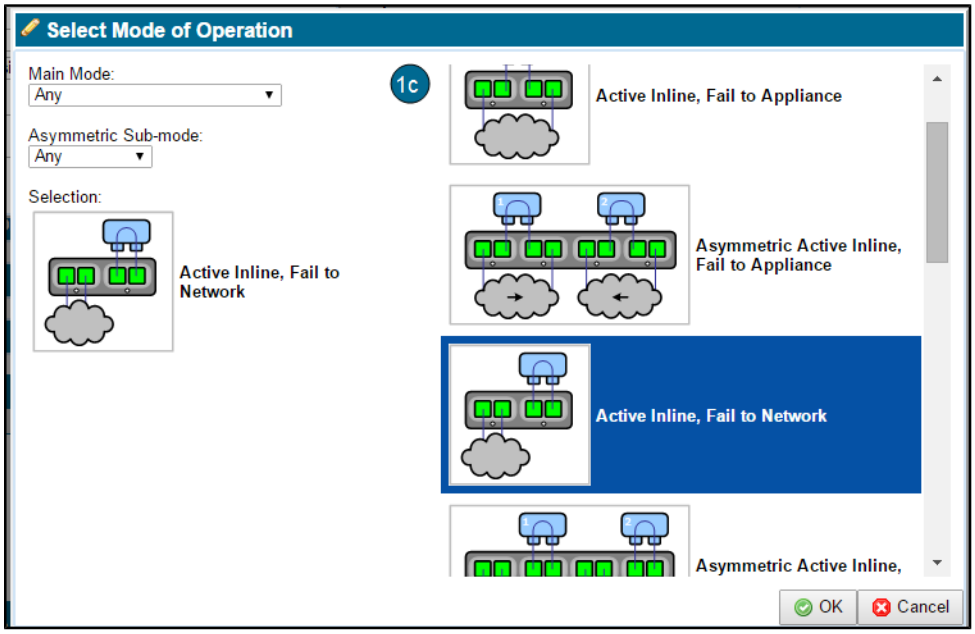

Step 2 On the **Add Segment** window, select a previously defined **Ruleset**. The specific choices available will depend on your segment and deployment mode.

Step 3 At **Session Log Mode**, select where to save log files.

Step 4 Select **VLAN Translation** to enable VLAN mapping between interfaces.

Step 5 Click **OK**. The window closes.

Step 6 To configure the VLAN mapping, edit the **VLAN Mappings** panel at the bottom of the window. Click **Add** in the **VLAN Mappings** panel header. The **Add VLAN Mapping** window displays.

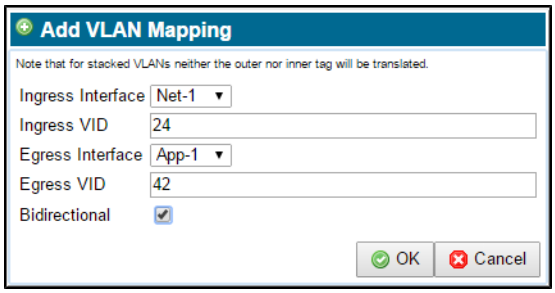

- Configure your VLAN translation to apply a new VLAN mapping to packets between the **Ingress Interface** (**Net-1** by default) port and **Egress Interface** (**App-1** by default).
- Supply the **Ingress** and **Egress VIDs** (VLAN identifier).
- Click **Bidirectional** to create a corresponding rule for the reverse direction. For example, if you have created a rule such that Net-1 with VID 24 maps to Net-2 with VID 42, then the reverse Net-2 with VID 42 maps to Net-1 with VID 24 rule is created.
- The rules are validated once they are configured. If there is an error in the configuration, you will see a warning message.

Step 7 Activate the segment.

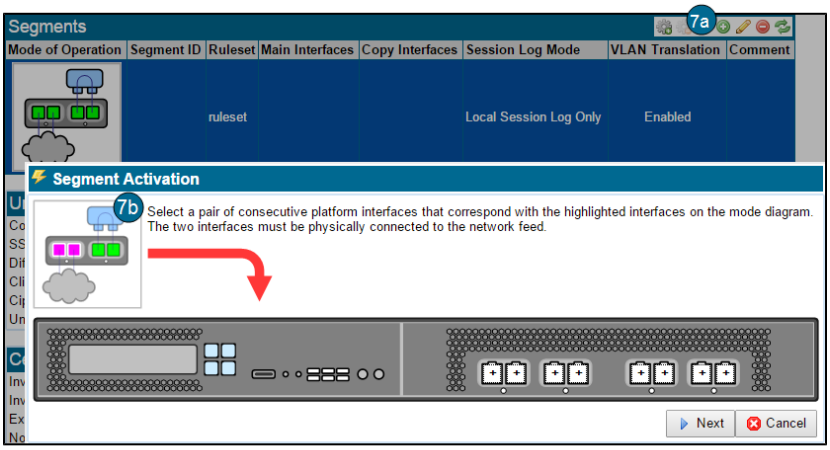
7a
Segments
Mode of Operation
Segment ID
Ruleset
Main Interfaces
Copy Interfaces
Session Log Mode
VLAN Translation
Comment
ruleset
Local Session Log Only
Enabled
Segment Activation
7b
Select a pair of consecutive platform interfaces that correspond with the highlighted interfaces on the mode diagram. The two interfaces must be physically connected to the network feed.
Next
Cancel

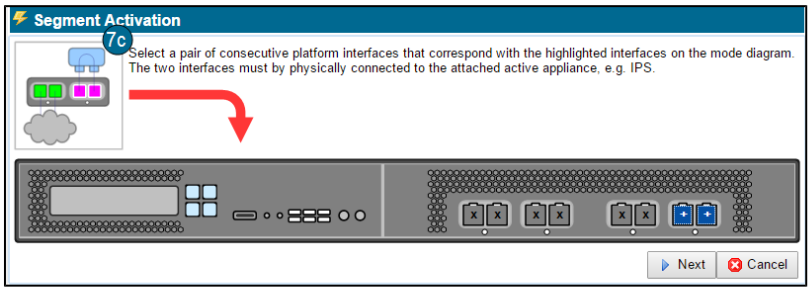
Segment Activation
7c
Select a pair of consecutive platform interfaces that correspond with the highlighted interfaces on the mode diagram. The two interfaces must by physically connected to the attached active appliance, e.g. IPS.
Next
Cancel

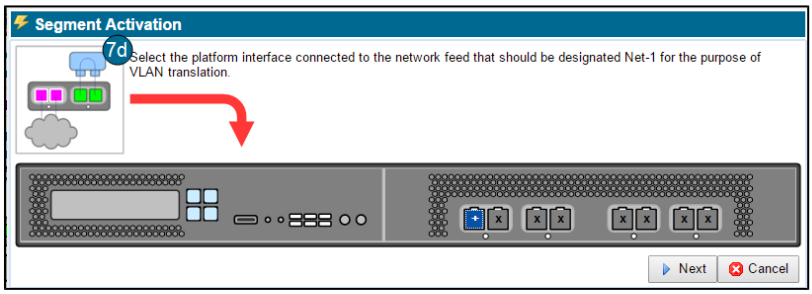
Segment Activation
7d
Select the platform interface connected to the network feed that should be designated Net-1 for the purpose of VLAN translation.
Next
Cancel

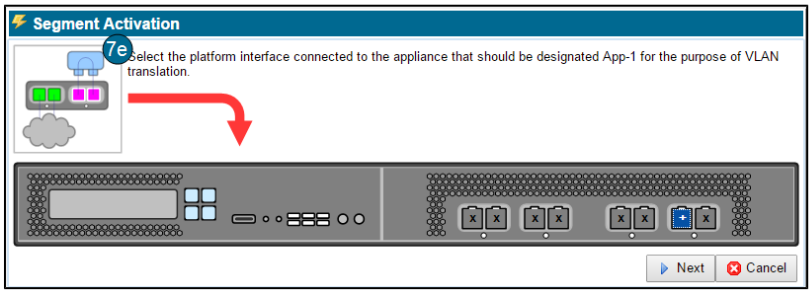
Segment Activation
7e
Select the platform interface connected to the appliance that should be designated App-1 for the purpose of VLAN translation.
Next
Cancel

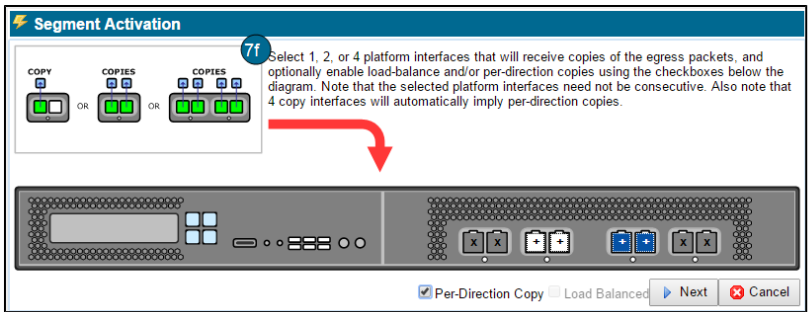
Segment Activation
7f
COPY
COPIES
COPIES
OR
OR
Select 1, 2, or 4 platform interfaces that will receive copies of the egress packets, and optionally enable load-balance and/or per-direction copies using the checkboxes below the diagram. Note that the selected platform interfaces need not be consecutive. Also note that 4 copy interfaces will automatically imply per-direction copies.
Per-Direction Copy
Load Balanced
Next
Cancel

- On the **Segments** window, click the **Mark for Activation** tool (cog with green plus). The first **Segment Activation** window displays. The graphic shows you which interfaces are available to select at each step of the process.
- For this example, next choose the two interfaces connected to the network, then click **Next**.
- Select the two interfaces connected to the attached active appliance, then click **Next**.
- Select the SSL Appliance interface connected to the network interface which will be designated **Net-1**. This interface corresponds to **Net-1** in the **VLAN Mappings** table. Click **Next.**
- Select the SSL Appliance interface connected to the attached appliance interface which will be designated **App-1**. This interface corresponds to **App-1** in the **VLAN Mappings** table. Click **Next.**
  - Optionally, select one, two, or four interfaces to receive copied packets, then click **Next**.
  - Click **Per-Direction copy** to have all traffic copied to each passive appliance.
  - Click **Load Balanced** to balance the copied traffic between the passive appliances.
  - When VLAN is in use, copies sent to copy ports have the same translation as the corresponding appliance port (for example, the same as App-1).
- Click **Next**. The **Segment Activation** window closes.

Step 8 Click **Apply**. The segment will apply VLAN translation.

Subject/Domain Names List

Entries in a **Subject/Domain Names List** are matched against the domain names and certificate subject of the SSL server for a session. The server Common Name (CN) and Subject Alternate Names (SAN) fields in the SSL server certificate are used in addition to the Server Name Indication (SNI) field from the ClientHello message.

The set of server domain names derived from the SSL handshake is used to match against the Subject/Domain name values specified in a rule, and if one of them matches, the rule will be triggered, and the appropriate policy applied. The server domain name displays in the SSL session log (SSL Session Log, page 6-10).

The **Subject/Domain Names List** display contains two panels. See the Overview of Common Tools, page 6-5 for information on using the tools. A **Subject/Domain Names List** called **sslng-unsupported-sites** is configured by default. It contains the domain names of SSL sites, the traffic to which cannot be inspected. Selecting the list in the upper panel causes the set of names in the list to display in the lower **Subject/Domain Names** panel. The figure shows the first page of names in the default **sslng-unsupported-sites** list.

Note A cut-through rule using the **sslng-unsupported-sites** list should be included in the ruleset used on any in-line segment in order to enable applications using these sites to function normally.

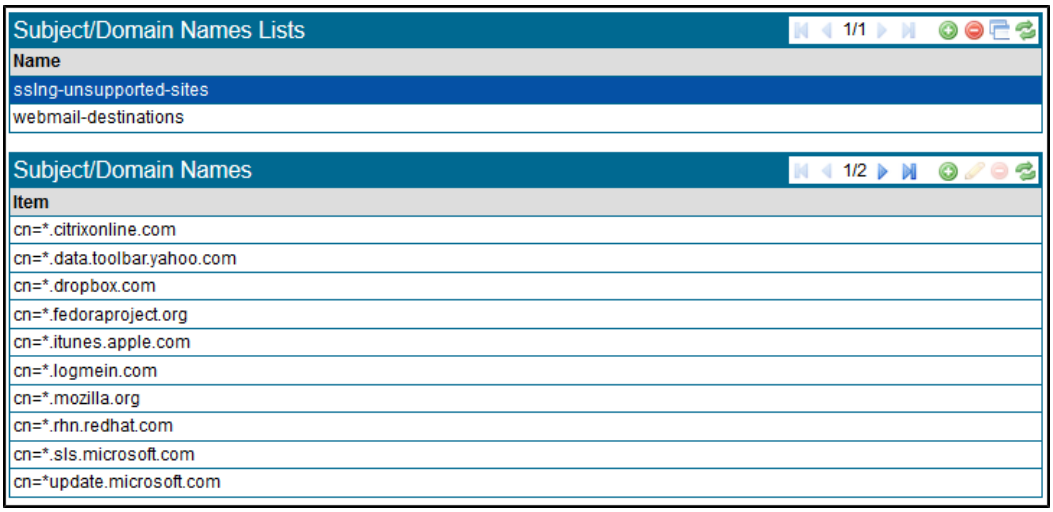

Click **Add** in the **Subject/Domain Names List** panel, to bring up the **Add Subject/Domain Names List**. Enter the name of the new list, then click **OK**.

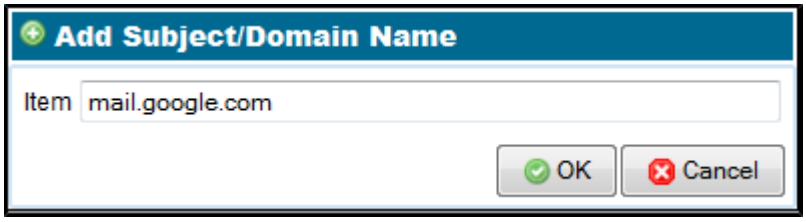

Domain Names entered here can begin with the wild card "*" character. For example, “*.example.com” will match flows to all example.com subdomains. Subject distinguished name attributes can be entered using CN=, O=, OU=, and C= DN attribute prefixes. The following example shows how a subject DN might be entered using this syntax:

- *cn= www.example.com
- CN=*.example.com, OU=Research, O=Example, Inc., C=US

The entries are case insensitive.

Domain Names List

Use **Domain Names Lists** to use a list of domain names as a rule match field. **Domain Names Lists** can only contain domain names, and not subject distinguished name attributes. When a **Domain Names List** rule match field is used, the SSL Appliance deduces the SSL flow domain name and compares it against the domain names in the list.

Searching **Domain Names Lists** is optimized, so these lists can contain many thousands of entries. A typical use for **Domain Names List** 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.

Maintaining large **Domain Names Lists** using the WebUI is a very manual task. External tools that simplify and automate the management of such lists might be available to simplify this task.

See the Overview of Common Tools, page 6-5 for information on using the tool icons. The **Remove** and **Clone** tools are inaccessible unless an entry in the table is selected.

The figure presents the Domains Names panels, including how to add a new Domain Name.

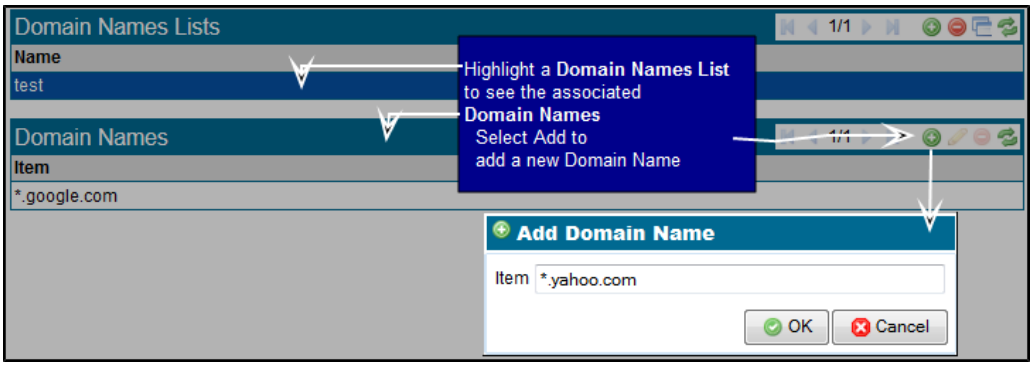

IP Address Lists

A typical use for an **IP Addresses List** is to prevent inspection of traffic to many different sites of a particular type based on the destination IP address of the hosts.

The **IP Addresses Lists** window contains two panels. The lower **IP Addresses** panel content varies depending on the row selected in the upper **IP Addresses Lists** panel. Each IP Addresses list occupies one row. Searching is optimized so that these lists can contain many thousands of entries.

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 these 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
- IPv6 addresses may be full or collapsed; see IP Address Format Notes, page 6-50 for more details.

Addresses are validated on input so the system will not allow input of an illegal IP address. See the Overview of Common Tools, page 6-5 for information on using the tool icons.

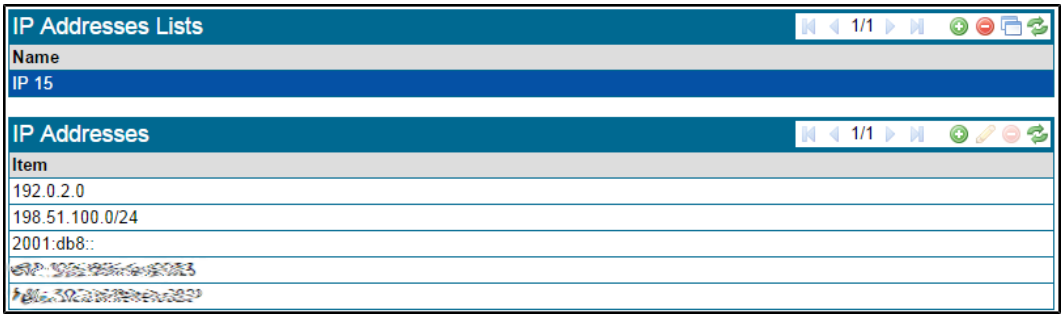

The figure shows the **IP Addresses** panel with several addresses entered, in IPv4 and in IPv6 collapsed and full formats. Maintaining large **IP Addresses Lists** using the WebUI is a very manual task. External tools that simplify and automate the management of such lists may be available to simplify this task.

Cipher Suites Lists

The **Cipher Suites Lists** window contains two panels, the upper **Cipher Suites Lists**, and the lower **Cipher Suites**. Select a list in the upper panel to display the set of cipher suites in the list in the lower panel. Each **Cipher Suites List** occupies one row. See the Overview of Common Tools, page 6-5 for information on using the tool icons.

When adding a cipher suite to a list, the **Add Cipher Suite** window displays. Select the additional cipher suite from a drop down list, or input it as a number in decimal or hex format

User-imported RSA key sizes are limited to 2048-bit and 3072-bit.

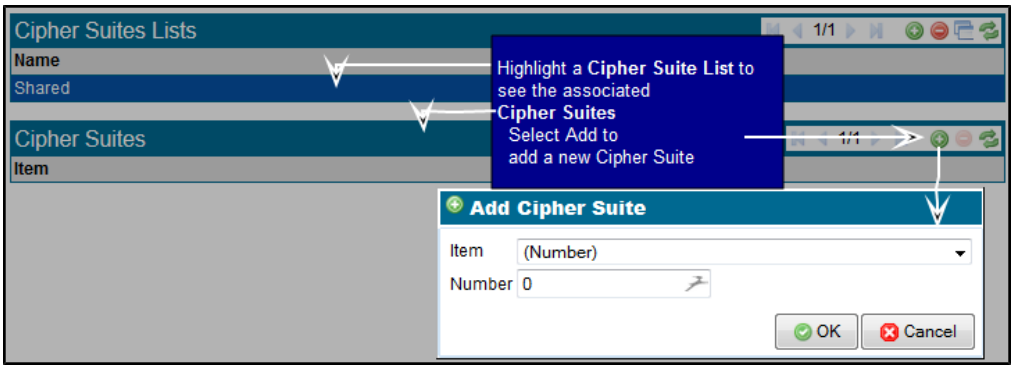

The next figure shows a list with three entries, each using a different input format. The drop down menu provides a list of all cipher suites using the name format, for example, TLS_RSA_SHA_WITH_AES_CBC_SHA.

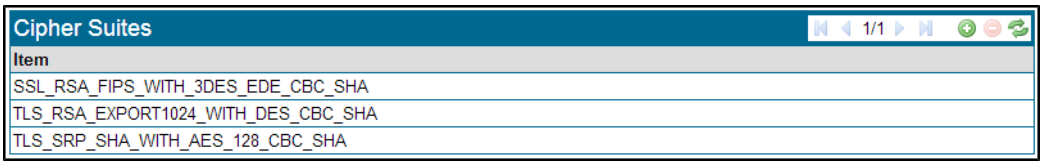

Host Categorization Lists

Use this window to view and manage Cisco Host Categories. The Cisco Host Categorization service allows policy to be tailored to the destination of an SSL flow. With this feature enabled, you can write policy specific to a type of traffic. For example, you could configure a policy to cut-through all traffic to financial services sites. See the Overview of Common Tools, page 6-5 for information on using the tool icons.

The SSL Appliance matches categories found in SSL flows and applies the policy. The updated database downloads periodically. The currently configured settings appear on initial view.

Notes

- The Cisco **Host Categorization** service requires a valid license. See License, page 6-69. It uses a database that must be downloaded from Cisco. Proper credentials are required to download the database.
- It might take up to six hours for the downloaded license credentials to be ready to use. If you can't download the database initially, wait up to six hours, and try again.

Use the **Host Categorization Status** area to get a snapshot of the current state of your Host Categorization database, with information such whether a download is in progress, and the state of the license.

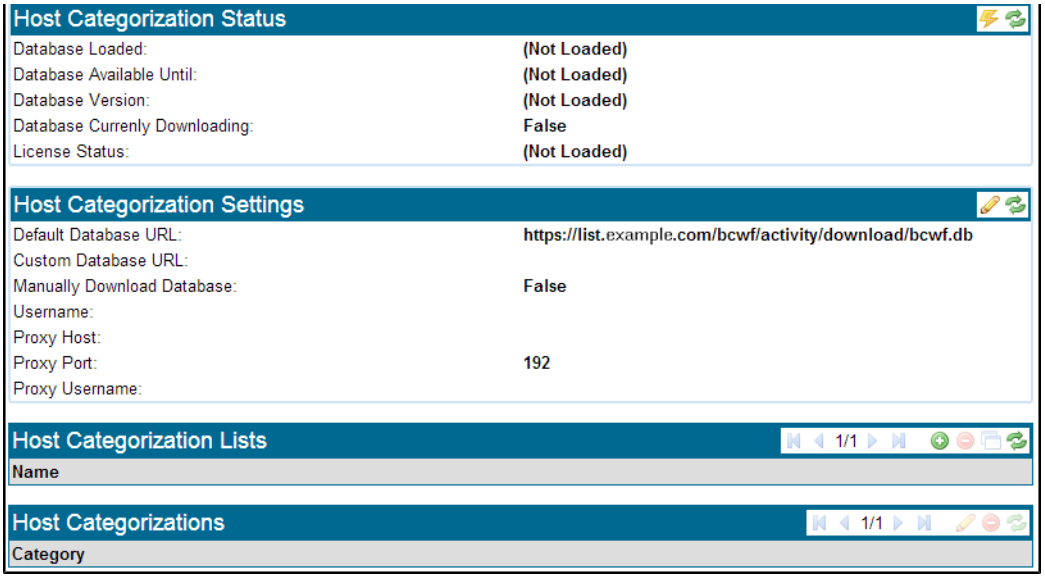

Download the Host Categorization Database

The Cisco Host Categorization service uses a database that must be downloaded from Cisco. The database is approximately 500 Mb in size and might take several minutes to download; it might temporarily use about 1G of space as it initializes. Use the **Host Categorization Status** panel to view and manage the database, and the **Host Categorization Settings** panel to view and manage the connection settings.

The first time you use the **Host Categorization List**, you must first download the Host Categorization database (license is required). When you update the download settings, the download begins automatically. If you have selected **Manually Download Database**, click the lightning bolt icon to start the download. You will see a confirmation message.

A **Database Currently Downloading: True** status message will appear in the Status window. Once installed, the database automatically updates every five minutes for the default URL (two hours if other), unless you have selected Manually Download Database.

Database Download Tips

Click **Add** only once.

- Refresh the window to see if the download has completed; the Database Loaded setting will indicate the download date, and the **Database Currently Downloading** status will read **False**.
- Click Apply to confirm your changes.
- Check the System Log (View System Log Entries, page 6-9) for warning messages.

To change the settings, click **Edit** in the far right of the **Host Categorization Settings** title bar. The **Edit Host Categorization Settings** window displays.

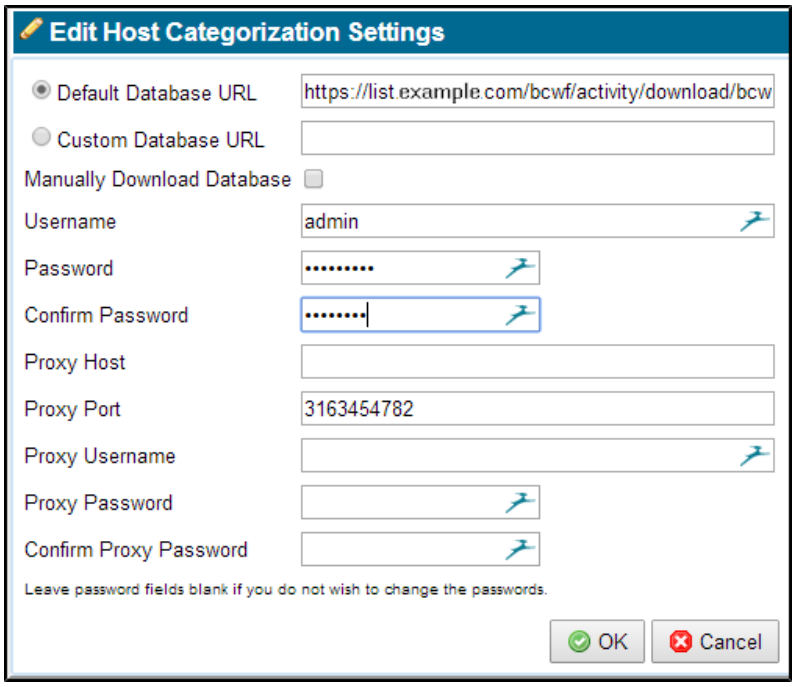

Usually, you will select the Default Database URL to use the Cisco supplied path to the categories database, and let it update automatically. After entering the Username and Password to download the database the first time, you don't need to enter that data again, unless you are changing the values. These settings apply to the download site not the SSL Appliance.

Use the Host Categorization Lists

Maintain or view your categorization lists in this panel.

The categories database (located at https://list.bluecoat.com/bcwf/activity/download/bcwf.db) may be downloaded securely through the SSL Appliance, downloaded to a local web server and applied from there, or downloaded through a proxy. To use a proxy, set the proxy host and port. If required, also set the proxy username and password.

Create a New Host Categorization List

Step 1 Click Add.

Step 2 Enter the list Name on the **Host Categorization List** pop up.

Step 3 Click OK.

To see what categories are included in a **Host Categorization List**, highlight the list name. The corresponding categories appear under **Host Categorizations**.

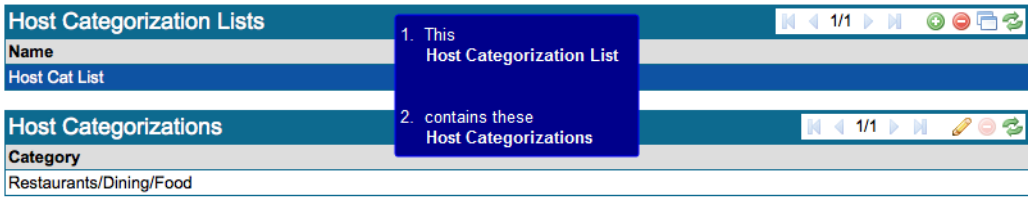

Add Categories to a List

Step 1 Highlight the row of the **Host Categorization List** you want to edit.

Step 2 Under **Host Categorization**s, click Edit. The **Change Selected Categories** window opens.

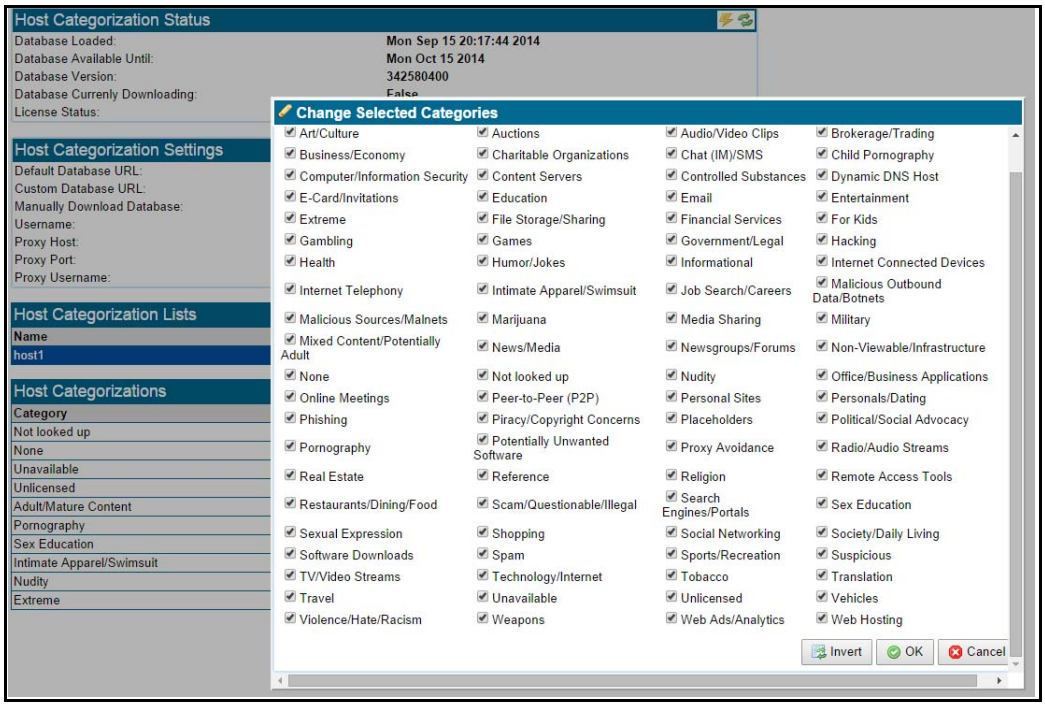

Step 3 Select the required categories. Click Invert to reverse your selection; for example, if you have selected all of the categories, clicking Invert will select no categories.

Step 4 Click OK.

Note The categories displayed might change, depending on the database.

Delete Categories from a List

Highlight the category under **Host Categorization**, and click **Delete**. Alternately, deselect the category in the **Change Selected Categories** window.

Examples of Category Usage in Policy

- Use rules in your policy (seeConfigure Segments and Policies, page 6-17) ruleset (seeConfigure Rulesets to Handle SSL Traffic, page 6-18) to match SSL flows to host categories.
- Create a rule which will cut-through traffic that matches the selected category list, and decrypts everything else.
- Create a rule where only traffic matching the list will be decrypted (everything else is cut-through).

Rename a Category

Category names may be removed, added, or changed when the database is updated, which can affect policy. Category renames are processed automatically, and a system log is generated if the rename results in a change in policy. Removed categories will be highlighted in red in the policy. A flow cannot match a removed category name.

System Log Data

The following Host Categorization licenses warnings and errors are reported in the System Log (View System Log Entries, page 6-9).

- An INFO message when the version of the database changes.
- WARNING message will be made 15, 5, 4, 3, 2, and 1 days before the database becomes stale.
- An ERROR message when the database becomes stale.
- A WARNING message will be made 30, 15, and 5 days before the Host Categorizations license expires.
- A WARNING level system log entry will be made every day during the last 5 days before the license expires.
- An ERROR level system log entry when the license expires.

If the database becomes stale, the flow will be categorized as "Unavailable."

A valid Cisco Host Categorization component license will be required to categorize flows. Without a license, flows will be categorized as “Unlicensed.”

Session Log Data

The **Session Logs** (SSL Session Log, page 6-10) include Host Categories information:

- The first specific Host Category matched by a flow (hence triggering a ruleset); only one category is included in the log, even if the flow matched multiple categories, and more than one category triggers the rule.
- The SNI for a session; this will help in troubleshooting Host Categorization issues, as you will be able to identify the site the user was trying to visit.

No **Host Categorization** information is included in the **Session Log** if no rule is matched. The Session Logs data can be exported for off-box analysis.

Traffic Classes Lists

Use **Traffic Classes Lists** to construct policy which decides whether or not to intercept an SSL flow based on QoS bytes. The SSL Appliance looks at the IPv4 Type of Service or IPv6 Traffic Class byte in ClientHello packets.

See the Overview of Common Tools, page 6-5 for information on using the tool icons. Highlight an entry in the **Traffic Classes Lists** panel to view related content in the **Traffic Classes** panel.

Configure Traffic Classes

Step 1 Create a new Traffic Class List; click Add in the upper **Traffic Classes Lists** panel. On the **Add Traffic Classes List** window, enter the name of the new list. Alternatively, you may edit a preexisting list.

- Lists are especially useful when you need to match more than one value to see the traffic of interest.

Step 2 Add a Traffic Class to a list: Highlight the **Traffic Classes List**, then click Add in the **Traffic Classes** panel. Enter the required **Value** and **Mask**.

- Values are displayed in hex, but may be entered in hex or decimal
- **Value** range: 0-255/0x00-0xFF
- **Mask** range: 0-255/0x00-0xFF; a **Mask** value of 0x00 will match all traffic classes, regardless of the **Value** entered
  The default **Mask** value of 0xFC matches DiffServ values
- Traffic Class matching is not limited to standard QoS values; you can specify any value and mask pair

Step 3 Click **OK** and **Apply** to save and implement the changes.

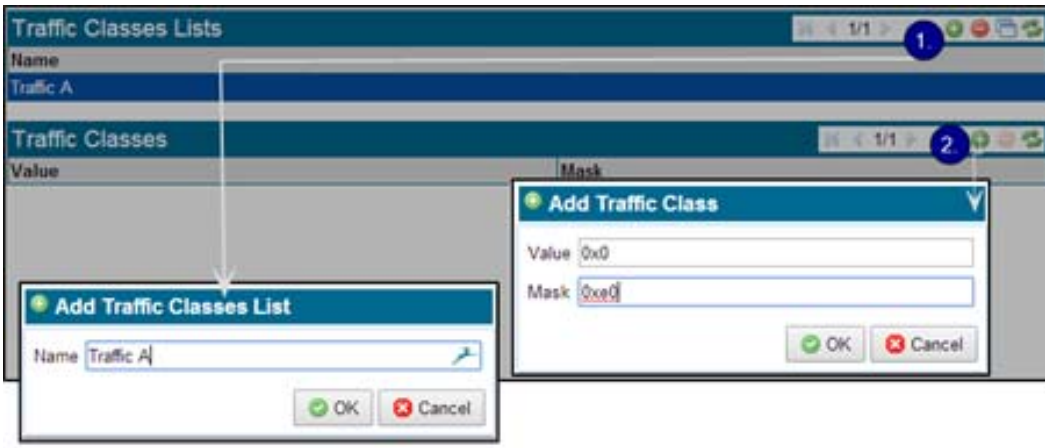

Use Traffic Classes in Policy

The **Policies > Rulesets > Insert Rule** and **Edit Rule** panel contains three Traffic Class related options, as follows. Traffic Class values may be set directly on the panel, or a ruleset may refer to a **Traffic Class List**.

Traffic Class Unconfigured	Policy will not be based on Traffic Classes
Traffic Class (Value)	Enter the specific **Value** and **Mask** values here; rules refers to these values only
Traffic Class List	Select a Traffic Class List (configured in (**Policies > Traffic Class Lists**) to use as reference for policy

Policy Examples

- You want to inspect "normal priority" IPv4 traffic, defined with the specific following values.
  - On the **Policies > Traffic Classes List** window, create a new **Traffic Classes List.**
  - Add two entries: check **Traffic Class Values**, and enter these values

	Rule 1	**Rule 2**
Value	0/0x00	96/0x60
Mask	224/0xE0	224/0xE0

  - Configure policy which uses this **Traffic Classes List**; on the **Policies > Rulesets > Insert Rule** panel, select **Traffic Classes List**, and choose the new list from the drop down. See the next figure.
  - Apply the changes.
- You want to match flows containing all Traffic Class values with a 0 in the least significant bit, and with a Traffic Class value of 16/0x10.

Step 1 Create a new **Traffic Classes List**.

Step 2 Add one **Traffic Class** to the new **Traffic Class List**, with **Value** = 0/0x00 and **Mask** = 1/0x01 (matches traffic with a 0 in the least significant bit).

Step 3 Add one **Traffic Class** with **Value** = 16/0x10 and **Mask** = 255/0xFF (matches Traffic Class value 16/0x10).

Step 4 Create a **Ruleset** with a rule referencing this **Traffic Class List** (tick **Traffic Class List**, and select the new list from the **Traffic Class List** options).

Step 5 Apply the changes.

PKI Management

The **PKI** menu contains options for managing certificates and keys, and for creating lists of certificates and keys. Each menu option is described below.

Note A user must have the Manage PKI role in order to make changes to the certificates and keys on the system. Users without the Manage PKI role will find that some features of the PKI menu are not available to them.

Resigning Certificate Authorities

Use the **Resigning Certificate Authorities** window to create, import, export and manage Certificate Authorities. The PKI store includes HSM resigning CAs, with private keys stored on the HSM appliance, as well as local resigning CAs, with private keys stored internally on the SSL Appliance.

Certificate Tools

	Add certificate.			Multipage tools
	Delete certificate.			View certificate details
	Edit			Generate certificate
	Export certificate			Refresh

Install a Local CA for Certificate Resign, page 4-8 describes the different ways a local certificate resigning CA can be added to the system. Multiple resigning Certificate Authorities can be configured and stored in the system. The choice of which resigning CA is used to resign a server certificate when an SSL session is being decrypted using certificate resign is controlled by either the segment, ruleset or rule definition. Which resigning CA is used can be configured to depend on details of the server certificate for the session being inspected, so different resigning CAs can be used for traffic going to different servers over the same segment.

External Certificate Authorities

The **External Certificate Authorities Lists** window contains two panels. Selecting a list in the upper **External Certificate Authorities Lists** panel causes the set of **External Certificate Authorities** certificates in the list to display in the lower panel. Each External Certificate Authorities list occupies one row in the **External Certificate Authorities Lists** panel. See Certificate Tools, page 6-43 for details on using the tools.

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. Selecting this list in the upper panel will cause the lower External Certificate Authorities panel to display details of the CA certificates in the list.

Click **Add** on the **External Certificate Authorities Lists** panel to create and add a custom list. Select the new list, then copy CA certificates from the **all-external-certificate-authorities list** to the new custom list.

The custom list is always a subset of the **all-external-certificate-authorities** list; it cannot contain entries that are not present in the **all-external-certificate-authorities** list. Select a custom list and click **Add** in the lower panel; a window displays where you can add keys in the default list to the custom list.

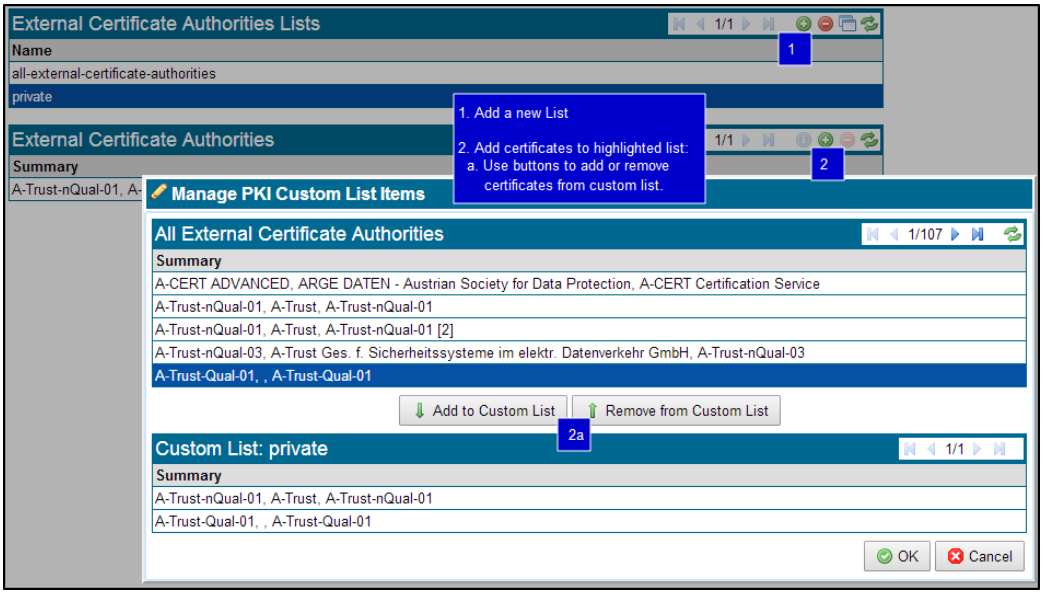

The figure shows an example where two CA certificates from the **all-external-certificate-authorities** list have been added to a custom list called “private”. One of the entries that has been included in the private list is a private CA certificate that had previously been imported to the **all-external-certificate-authorities list:** the Cisco Systems CA.

Use the **Clone** tool on the **External Certificate Authorities Lists** panel 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 certificates to or from the new version than to start from scratch.

Certificate Revocation Lists

The **List of Certificate Revocation Lists** display contains two panels. See Certificate Tools, page 6-43 for details on using the tools.

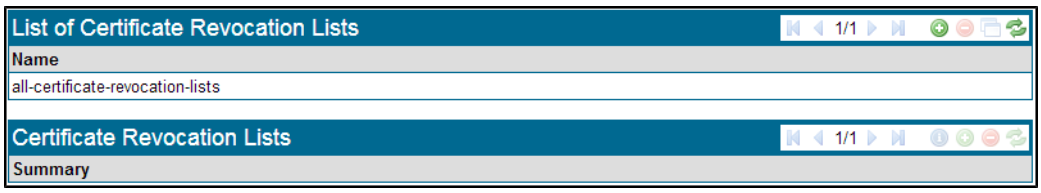

When populated, selecting a list in the upper panel causes the set of CRLs in the list to be displayed in the lower panel. Each **Certificate Revocation List** occupies one row in the **List of Certificate Revocation Lists** panel.

The system has a default list installed, the **all-certificate-revocation-lists** list. This list is initially empty. Select this list to see a display of the CRLs in the list in the **Certificate Revocation Lists** panel. Select this list and click **Add** in the lower **Certificate Revocation List** panel to open up a window where you can import a CRL.

If the CRL file being imported is encrypted and protected with a password, in the **Import CRL List** window you must enter the password in the **Password** field.

Create a Custom Certificate Revocation List

Click **Add** on the **List of Certificate Revocation Lists** panel to create and add a custom list. Once this list is created, select it, and copy CRLs from the **all-certificate-revocation-lists** 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. When a custom list is selected and **Add** in the **Certificate Revocation Lists** panel is clicked, a window displays where you can add keys in the default list to the custom list.

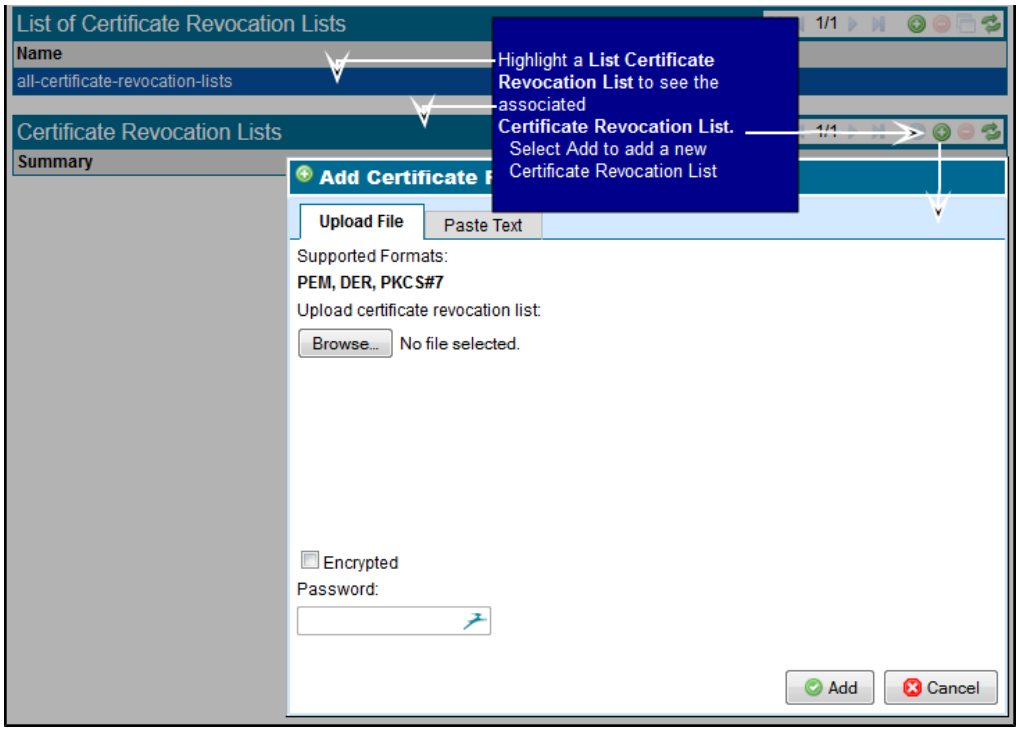

Cloning a Certificate Revocation List

Use the clone feature on the **List of Certificate Revocation Lists** panel 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 CRLs to the new version produced by the **Clone** tool.

Trusted Certificates

The **Trusted Certificates** display contains two panels. See Certificate Tools, page 6-43 for details on using the tools. Select a list in the upper panel to view the set of certificates in the list in the lower panel. Each **Trusted Certificates** occupies one row in the **Trusted Certificates Lists** 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 list in the lower **Trusted Certificates** panel. Select this list and then click **Add** in the lower **Trusted Certificates** panel to open up a window where you can import a certificate.

Click **Add** on the **Trusted Certificates Lists** panel to create and add a custom list. Once this list is created, select it, and then copy certificates from the **all-trusted-certificates** list to the custom list as required.

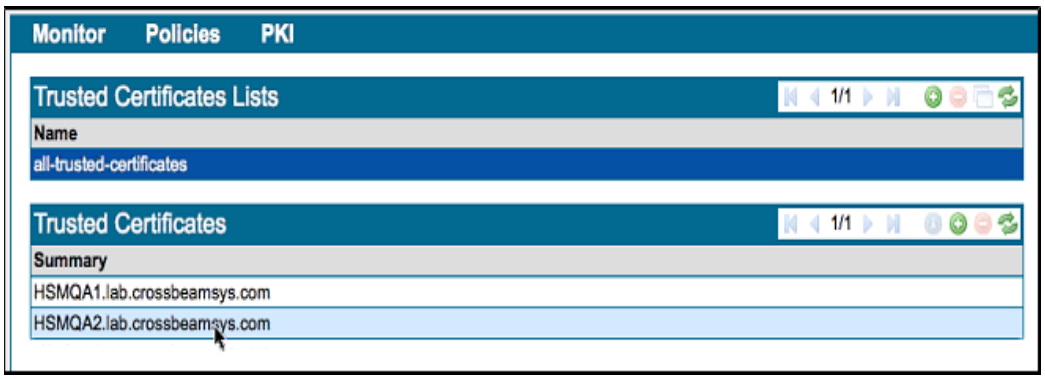

The custom list is always a subset of the **all-trusted-certificates** list, and cannot contain entries that are not present in the **all-trusted-certificates** list. When a custom list is selected and you click **Add** in the lower panel, a window displays where you can add keys in the default list to the custom list.

Use the **Clone** tool on the **Trusted Certificates Lists** panel 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 than to start from scratch.

Known Certificates and Keys

The **Known Certificates and Keys** window contains two panels. See Certificate Tools, page 6-43 for details on using the tools. Selecting a list in the upper panel causes the set of certificates with keys in the list to be displayed in the lower panel. Each **Known Certificates and Keys** occupies one row in the **Known Certificates and Keys Lists** panel.

Tip It is often quicker to clone an existing custom list and then add or remove certificates to the new version produced (saved and renamed) by the **Clone** tool.

The system has a default list installed, the **all-known-certificates-with-keys** list. This list is initially empty. Selecting this list in the upper panel will cause the lower **Known Certificates and Keys** panel to display details of the certificates with keys in the list.

Add a New Certificate with Key

Step 1 Click Add in the **Known Certificates and Keys** panel. The **Add Known Certificate with Key** window displays.

Step 2 Install the certificate and key by one of these methods, after entering any required Password and selecting Encrypted if necessary:

- On the **Upload File** tab, click **Choose File** at both the **Upload Certificate** and **Upload key** areas to browse to the license file location and select it, then click **Add** at the bottom of the window. OR
- On the **Paste Text** tab, paste in previously copied text of the certificate and the key into the respective fields, then click **Add**.

Create or Manage a Custom Certificate with Keys 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.

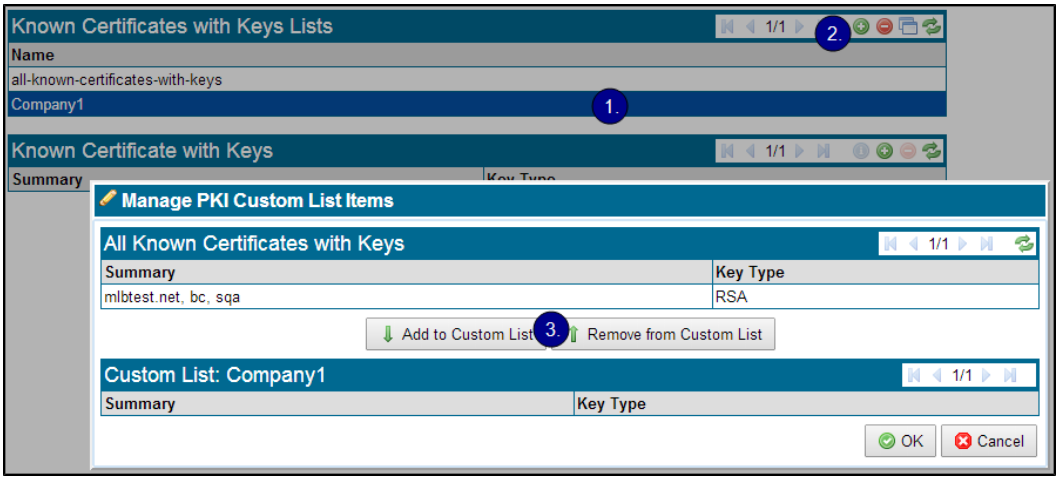

Step 1 Highlight a custom list in the **Known Certificates and Keys List** panel.

Step 2 Click **Add** in the **Known Certificates and Keys** panel. The **Manage PKI Custom List Items** window displays.

Step 3 Click **Add to Custom List** and **Remove from Custom List** to copy a known certificate with key to, or remove it from, your custom list.

Client Certificates

The SSL Appliance uses **Client Certificates** to mutually authenticate with an HSM. It functions in the same way as the other lists windows.

The RSA key size for generating client certificates and keys is 2048-bit

See Add Resigning Certificates, page 5-5 for details on using this window.

HSM Appliances

Use the **HSM Appliances** window to manage attached HSM appliance connections.

See Add an HSM, page 5-5 for details on using this window.

Platform Management

The **Platform Management** menu is to the right in the menu bar, and titled with the current hostname of the SSL Appliance. This menu includes tools for viewing and managing the platform, and for configuring and managing access to the platform network management features, which are described in the following sections. Platform management also includes managing user accounts and performing updates to the system software.

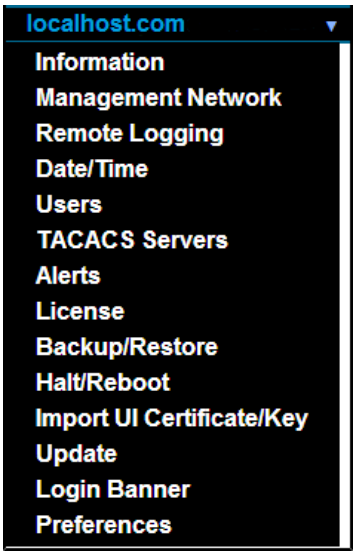

Information

View information about the software and hardware. The **Information** window initially shows two panels. The two panels have the **Refresh** tool for providing visibility of data, but no ability to enter or change data. Click **Show Advanced** to access additional information.

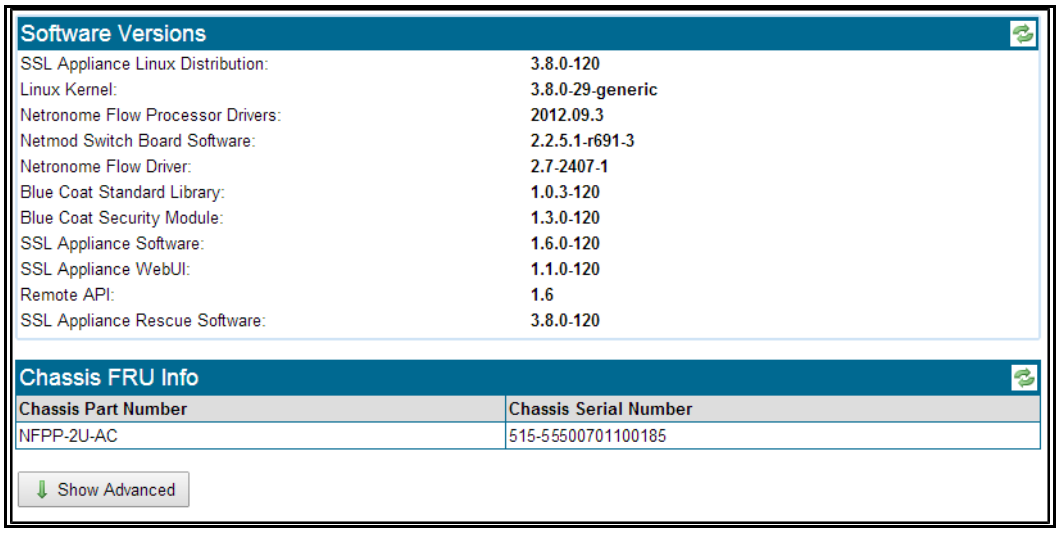

The upper **Software Versions** panel provides details of the software versions of the various software modules within the system. The **SSL Appliance Linux Distribution** value, in this example 3.8.3-118, is the most important element here, as this is the version number of the software that is running on the system. Cisco personnel might request the details from this panel when providing support for the device. Providing these details when filing a support ticket is useful.

The **Chassis FRU Info** displays in the lower section. Cisco personnel may request the details from this panel when providing support for the device. Providing these details when filing a support ticket is useful.

Click **Show Advanced** to see additional display-only information. These panels provide data on different hardware elements of the system. Cisco personnel might 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 card(s) installed in the system
- **BIOS and BMC Version**: BIOS details

Note The BIOS and BMC version data is also visible on the LCD screen.

Management Network

Use the **Management Network** window to configure the management IP, and Access Control Lists.

Management Network Panel

View and edit basic management network details. See the Overview of Common Tools, page 6-5 for information on using the tool icons.

The figure shows the **Edit Management Network** window used to adjust the network settings.

Note Set the **Hostname** of the localhost to its DNS hostname for the appliance.

Click **OK** and **Apply** as required to save your changes.

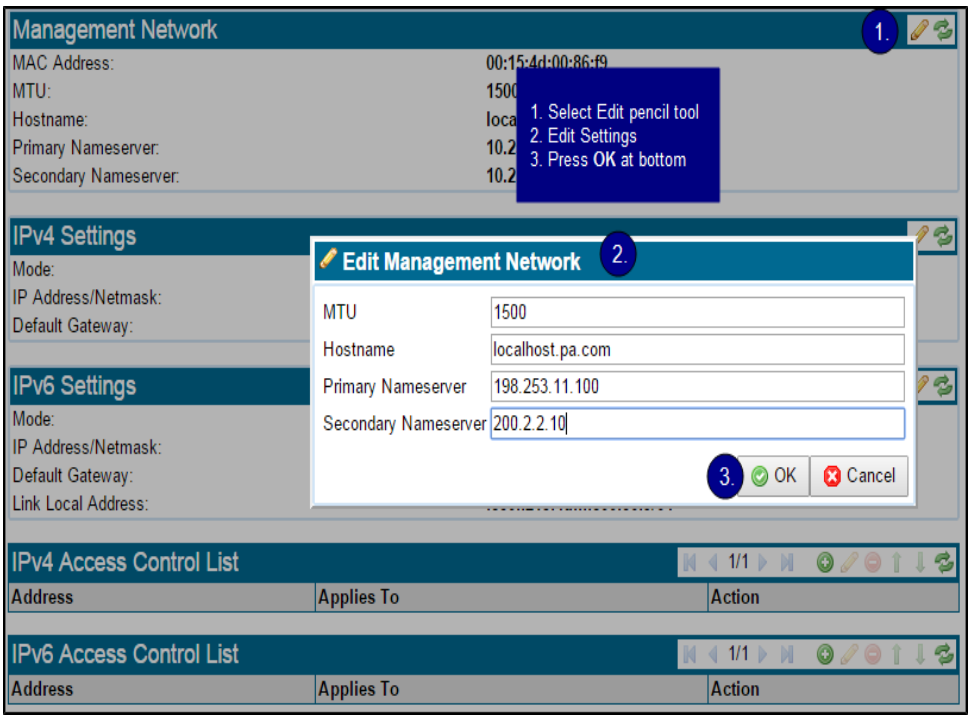

Note The SSL Appliance accepts both IPv4 and IPv6 configurations for IP Address configuration.

IPv4 and IPv6 Settings Panels

For IPv4 or IPv6 management, the appliance can be configured by the **Mode** setting to use either a **Static** (fixed) IP address or to acquire an IP address automatically using **DHCP** or a **SLAAC** method (IPv6 only). For DHCP to work there must be a working DHCP server on the network that the management Ethernet is connected to. DHCP is the default. See Configure a Static IP Address, page 4-4 for details on configuring a static IP address.

IP Address Format Notes

- Use the IP address/mask bits (CIDR) format to enter the IP address and netmask for all IP addresses.

  IPv4 Example: 192.0.2.0/24.

  IPv6 Example: 2001:db8::/24
- IPv6 addresses may be entered in full or collapsed form:

  Full: 2001:db8:0000:0000:0202:B3FF:FE1E:8329

  Collapsed: 2001::db8:0202:B3FF:FE1E:8329

IPv4 Settings Panel

Here is an overview of all of the IPv4 options.

Table 6-1 IPv4 Options

Mode	Setting	
	IP Address/Netmask	Gateway
Static	Required	Optional
DHCP	Automatically assigned (Not editable)	Automatically assigned (Not editable)
Disabled	Disabled	Disabled

Click Edit to change the settings. **IP Address/Netmask** and **Default Gateway** may be edited only when **Mode** is set to **Static**. Click OK and Apply as required to save your changes. See Configure Management Network Settings, page 4-3 for more details.

IPv6 Settings Panel

Here is an overview of all of the IPv6 options.

Table 6-2 IPv6 Options

Mode	Setting	
	IP Address/Netmask	Gateway
Static	Required	Optional
DHCP	Automatically assigned (Not editable)	Automatically assigned (Not editable)
SLAAC	Automatically assigned (Not editable)	Automatically assigned (Not editable)
SLAAC + Stateless DHCP	Automatically assigned (Not editable)	Automatically assigned (Not editable)
Disabled	Disabled	Disabled

Click **Edit** to change the settings.

- **IP Address/Netmask** and **Default Gateway** may be edited only when **Mode** is set to **Static**.
- When the **Mode** is set to **DHCP, SLAAC or SLAAC+Stateless DHCP,** that information is obtained automatically; choose the method appropriate to your network.
- The **IPv6 Link Local** address is derived automatically and presented; you may not edit it.

Click **OK** and **Apply** as required to save your changes. See Configure Management Network Settings, page 4-3for more details.

IPv4 and IPv6 Access Control List Panels

Access Control Lists (ACLs) are configured to authorize or restrict access to the management Ethernet interface (also called the “management network”). Independent ACLs are available for IPv4 and IPv6 traffic; you can also use an ACL to disable all SNMP access.. The default ACLs have no entries, and allows all access. To prevent access, traffic must match a **Block** entry before it matches an **Allow** entry. Traffic is allowed if it matches no entries in the ACL.

Access Control Lists apply to incoming connections only. See the Overview of Common Tools, page 6-5 for information on using the tool icons.

Note Configuring an **Access Control List** might terminate the management connection.

Configure an ACL List

Access Control List lists can have up to 1000 entries, but performance might suffer if a list includes over 200 entries.

- An ACL list entry may specify a single IP address, an IP subnet, or apply to any IP address. For example:

  IP: Enter IP address

  Subnet: Enter IP/mask

  All: IPv4 - 0.0.0.0/o; IPv6 - ::/0

  See the IP Address Format Notes, page 6-50 for details on entering IP addresses.
- Entries can apply to SSH/HTTP/HTTPS, ICMP/Traceroute, or SNMP traffic.

By default, connections are allowed. **Tip**: Use a catch-all entry at the end of the list to specify to **Deny** all connections; an **Allow** catch-all isn't necessary.

The first entry matched in a list ends the matching process.

- To configure a whitelist (connections accepted), set up an **Allow** entry on a list.
- To configure a blacklist (disallowed connections), set up a **Deny** entry on a list.

Use this process to configure an **Access Control List**, and refer to the figure following it.

Step 1 On the **(Platform Management) > Management Network** window, click Edit in the required **IPvX Access Control List** panel header.

Step 2 On the **Edit IPvX Access Control Rule** pop-up,

- Enter the incoming connection **Address**.
- Select the management category the ACL **Applies To**:
  - **SNMP**: SNMP traffic
  - **Network Utilities**: ICMP Ping, Traceroute
  - **Management** SSH/HTTP/HTTPS
- For the **Action**, select:
  - **Block** to reject a connection
  - **Allow** to accept the connection

Step 3 Click OK, then click Apply at the "Platform Config Changes" message.

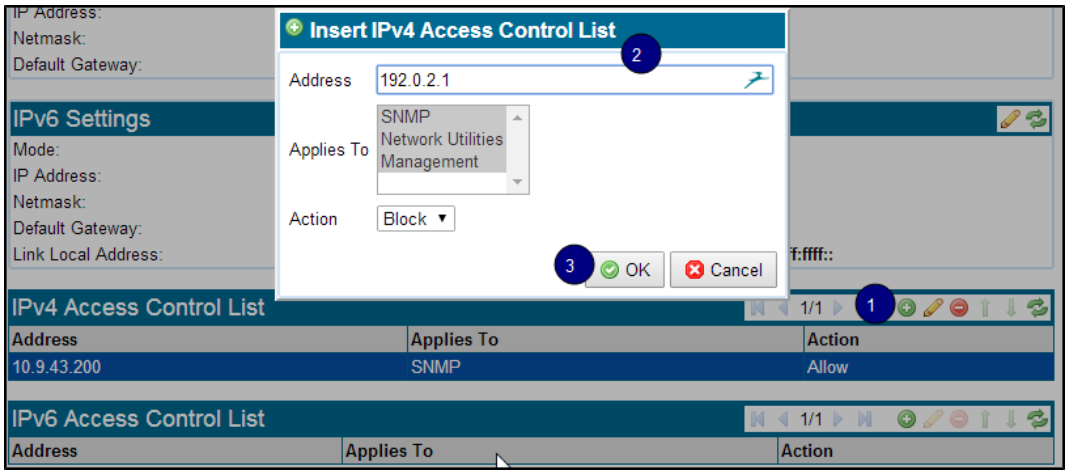

Troubleshooting

If you accidentally lock out the management interface due to misconfiguring the **Access Control List**, use the Command Line Diagnostics interface (see Command Line Diagnostics Interface, page 7-8) to edit the ACL.

Configure SNMP Access

Use the **(Platform Management) > SNMP Access** window to configure SNMP connections. SNMP is a standard protocol that allows access to the SSL Appliance for monitoring state and statistics through Get requests. Set requests can be used to write configuration, although their use is limited in the SSL Appliance to setting system parameters. Additionally, SNMP traps allow the appliance to notify a remote listener of state changes.

You can configure, enable, or disable SNMP management access; v1/2c and v3 may be enabled or disabled independently. For example, if you are using SNMPv3, you might want to disable SNMP v1/2c.

Starting at the top of this screen, optionally configure the **SNMP System Group** (Configure the SNMP System, page 6-54), then the required SNMP access: **SNMP v1/v2c Access** (Configure SNMP v1 and v2c Access, page 6-56) and/or **SNMP v3 Access** (About SNMP v3, page 6-57).

See the Overview of Common Tools, page 6-5 for information on using the tool icons.

Notes

- Brief events might not be reported through the SNMP MIB, or result in SNMP traps, based on limitations of the sensor sampling infrastructure and polling frequency.
- Click **Apply** after making any SNMP changes.
- Some SNMP configuration changes, such as changing a community string or a trap host, restart the SNMP agent and reset the SNMP sysUpTime when applied. SNMP sysUpTime is not related to the **Appliance Uptime** that displays on the **Dashboard** in the WebUI.

Get the MIBs

The SSL Appliance supports the standard SNMP MIB2 tables as well as five Blue Coat private MIBs:

- BLUECOAT -MIB
- BLUECOAT-LICENSE-MIB
- BLUECOAT-SEGMENT-MIB
- BLUECOAT-SG-SENSOR-MIB
- BLUECOAT-SG-USAGE-MIB

Blue Coat MIB MIB files are posted with the image file on BlueTouchOnline (https://bto.bluecoat.com/). The single zipped file contains the Blue Coat MIBs, and a reference to the supporting standard MIBs.

You must be logged in to access the files. If you don't have an account, go to the BlueTouch Request Login screen at https://www.bluecoat.com/forms/contact, and follow the registration process.

Download the MIBs

Step 1 Go to https://bto.bluecoat.com.

Step 2 Log in. The **Downloads** tab will not display unless you are logged in.

Step 3 Select **Downloads**.

Step 4 Select **Blue Coat Product Downloads**.

Step 5 Select your product.

Step 6 Select your appliance model (if applicable).

Step 7 Select a software version.

Step 8 Accept the Software Terms and Conditions.

Step 9 Download the MIB by clicking the file name links.

Note To load the Blue Coat MIBs on an SNMP network management station, place the Blue Coat and dependent MIBS into the folder designated by your SNMP manager tool, or use the SNMP manager's MIB import function.

Configure the SNMP System

You may optionally edit the values for certain MIB-2 system group objects, per RFC-3418.

Step 1 In the **SNMP System Group** header, click **Edit**. The **Edit SNMP System Group** window opens.

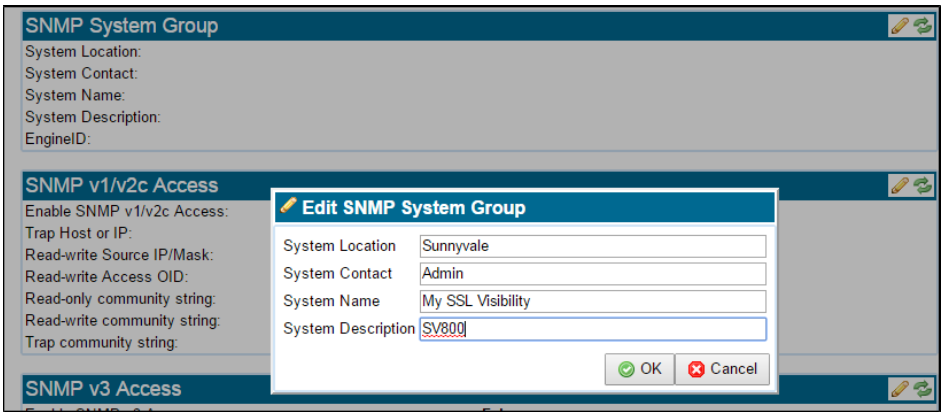

Step 2 Edit the system group information as required.

EngineID Notes

- The **EngineID** is not part of the MIB **SNMP System Group** and is read-only.
- Only SNMPv3 uses the **EngineID**, which must be unique among SNMP agents and systems that are expected to work together. The SSLV Appliance generates the value. This value persists across reboots.
- The **EngineID** value displays after SNMP v3 has been enabled and the change applied.It might take several seconds for the value to display.

Step 3 Click **OK**. The window closes.

Step 4 Click **Apply**.

Configure SNMP v1 and v2c Access

Step 1 Click **Edit** in the **SNMP v1/v2c Access** panel header to open the **Edit SNMP v1/v2c Access** window.

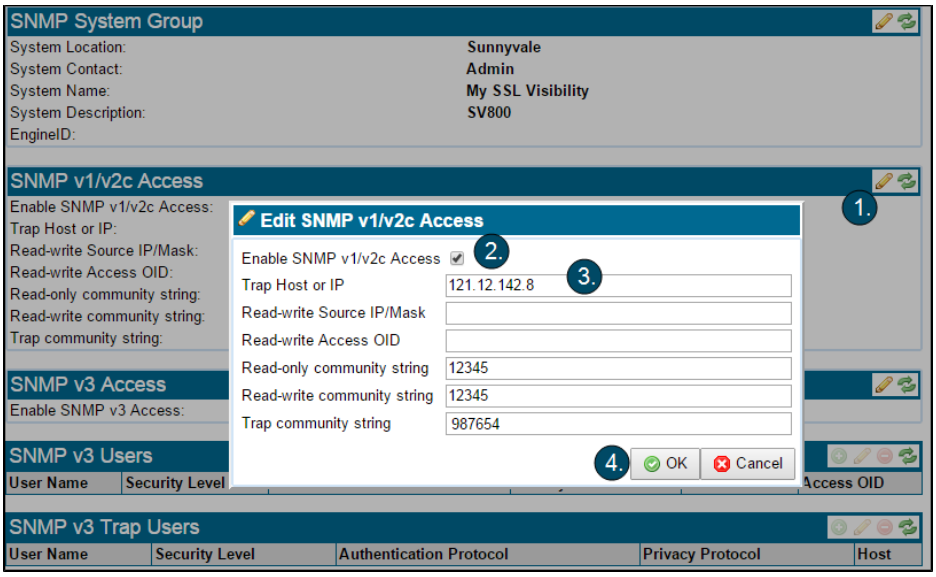

Step 2 Select **Enable SNMP v1/v2c Access** on the **Edit SNMP v1/v2c Access** window.

Note SNMP v1/v2c traps are disabled if **Enable SNMP v1/v2c Access** is not selected.

Step 3 Configure the parameters. At least one read-only or read-write community string is required. Community string values will appear on the WebUI

- **Trap Host or IP**: Enter the receiver (trap destination) IP address or hostname.
- **Read-write Source IP/Mask**: Optional: To restrict SNMP v1/2c read-write access to a single management station, enter the default source IP address for the MIB requester. Does not restrict read-only community requests.
- **Read-write Access OID**: Optional: To restrict SNMP v1/2c read-write access to a specific MIB subtree under the OID, enter the top level OID in dotted format (for example, 1.3.6.1.2.1.1). Does not restrict read-only community requests.

Note The **Read-write Access OID** configuration is only accepted when the **Read-write Source IP/Mask** is also configured.

- **Read-only community string:** Enter the value used to authenticate incoming SNMP read requests.
- **Read-write community string:** Enter the value used to authenticate incoming SNMP read-write requests.
- **Trap community string:** Enter the value used to authenticate outgoing trap messages.

Step 4 Click **OK**. The window closes.

Step 5 Click **Apply**.

About SNMP v3

SNMP version 3 provides for user authentication, encryption, and access control, as well as users for trap generation. Hence, you must create users when configuring SNMP v3. Up to 50 SNMP v3 users may be created. Once you have enabled SNMPv3 and created a credentialed user, you can access the appliance using SNMPv3. To generate traps, you must create a **Trap User** identity (see <LI Link text>"Create a SNMP v3 Trap User"). Access users and trap users must be unique.

Note Make sure to click **Apply** at the bottom of the **SNMP Access** screen when you make any SNMP changes.

Enable SNMPv3 Access

Step 1 Click **Edit** at **Enable SNMP v3 Access**. The **Edit SNMP v3 Access** window displays.

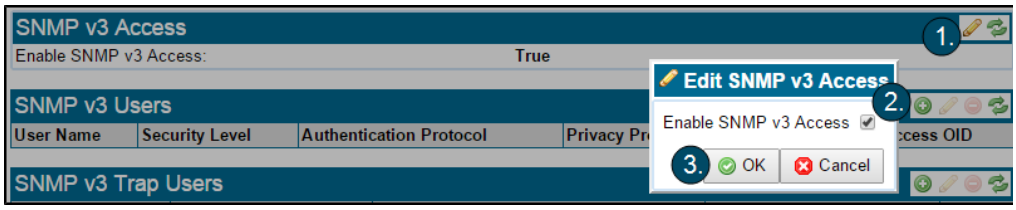

Step 2 Select **Enable SNMP v3 Access** on the **Edit SNMP v3 Access** window.

Note SNMP v3 traps are disabled if **Enable SNMP v3 Access** is not selected.

Step 3 Finish by clicking **OK**. The window closes.

Step 4 Click **Apply**.

Create a SNMP v3 User

An SNMP v3 user provides local access to SNMP from a remote SNMP query. For security purposes, passphrases once configured are obscured.

Step 1 Enable SNMPv3 access (see Enable SNMPv3 Access, page 6-57).

Step 2 In the **SNMP v3 Users** header, click **Add**. The **Add SNMP v3 Users** window displays.

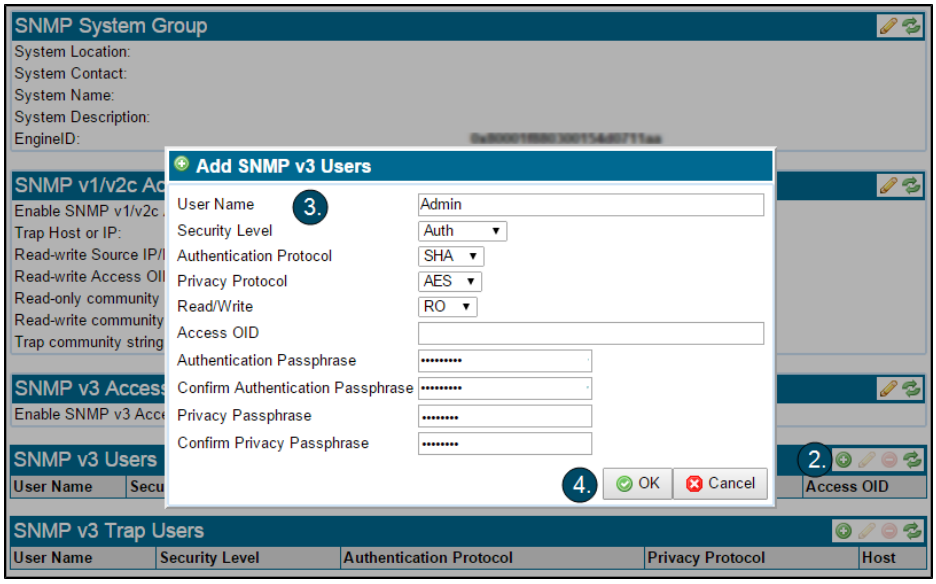

Step 3 Enter the user information:

- **User Name**: Must be between 4 and 31 characters in length, have no spaces, and not duplicate a **Trap User Name.**
- **Security Level**: Default **AuthPriv**; authentication and privacy required. **Auth** and **NoAuth** are also available.
- **Authentication Protocol**: Default **SHA**; this is the authentication algorithm HMAC-SHA-96. **MD5** (HMAC-MD5-96) is also available.
- **Privacy Protocol**: Default **AES;** this is the encryption standard CFB128-AES-128. **DES** (CBC-DES) is also available.
- **Read/Write:** Sets the access to MIB objects; default is **RO** (Read Only). **RW** (Read/Write) is available. The Blue Coat MIBs support only read-only elements. If the **System Group** objects are not configured with the WebUI, they are writable.
- **Access OID**: For View based SNMP v3 access, enter the top level OID in dotted format (for example, 1.3.6.1.2.1.1). It restricts access for the user to a MIB subtree under the OID specified.
- **Authentication Passphrase**: Required; between 8 and 31 characters with no spaces.
- **Confirm Authentication Passphrase**: Verify the **Authentication Passphrase** entry.
- **Privacy Passphrase**: Required; between 8 and 31 characters with no spaces.
- **Confirm Privacy Passphrase**: Verify the **Privacy Passphrase** entry.

Step 4 Click **OK.** The window closes.

Step 5 Click **Apply**.

Once a user exists, you can highlight the existing user and click **Edit** to edit the settings, or click **Delete** to remove that user.

Create a SNMP v3 Trap User

Configure **Trap Users** if you are sending traps. The **SNMP v3 Trap User** provides the user name used by the remote trap listener for traps sent by the local appliance. SNMP v3 supports the configuration of multiple trap destination users.

Follow this procedure to create an SNMPv3 **Trap User**. For security purposes, passphrases once configured are obscured.

Note Use the **EngineID** displayed in the **SNMP System Group** panel at the top of the window (see Configure the SNMP System, page 6-54) to configure the remote user to listen for traps sent from the appliance.

Step 1 Enable SNMPv3 access (see Enable SNMPv3 Access, page 6-57).

Step 2 In the **SNMP v3 Trap Users** header, click **Add**. The **Add SNMP v3 Trap Users** window displays.

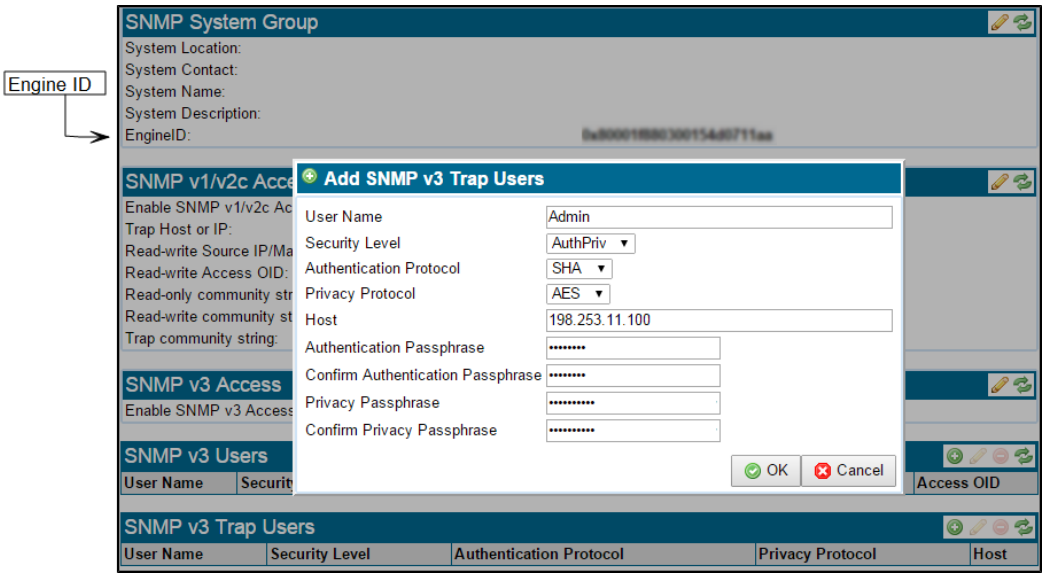

Step 3 Enter the new user information:

- **User Name**: Must be between 4 and 31 characters in length, have no spaces, and not duplicate an **Access User** name.
- **Security Level**: Default **AuthPriv**; authentication and privacy required. **Auth** and **NoAuth** are also available.
- **Authentication Protocol**: Default: **SHA**; this is the authentication algorithm HMAC-SHA-96. **MD5** (HMAC-MD5-96) is also available.
- **Privacy Protocol**: Default **AES;** this is the encryption standard CFB128-AES-128. **DES** (CBC-DES) is also available.
- **Host**: Enter the receiver (trap destination) IP address or hostname.
- **Authentication Passphrase**: Required; between 8 and 31 characters with no spaces.
- **Confirm Authentication Passphrase**: Verify the passphrase entry.
- **Privacy Passphrase**: Required; between 8 and 31 characters with no spaces.
- **Confirm Privacy Passphrase**: Verify the passphrase entry.

Step 4 Click **OK**. The window closes.

Step 5 Click **Apply**.

Once a user exists, you can highlight the existing user and click **Edit** to edit the settings, or click **Delete** to remove that user.

SSL Appliance SNMP Traps

SNMP traps are generated in the event of a status change in non-management network interfaces, licenses, segments, power supplies, fan and temperature sensors, and resource utilization (CPU, Memory, and Disk). The following table presents the trap types supported on the SSL Appliances, including status codes (if applicable), and any additional values from the MIBs that are sent with the trap.

Trap	Status Codes	MIB Values Sent with Trap
LinkUp (Non-management interfaces)		IfIndex
LinkDown (Non-management interfaces)		IfIndex
Licenses Status Change	Active=1 Expired=2	ComponentName ExpireType ExpireDate
Segment Status Change	OK = 1 Failure = bitwise OR of SoftwareFailure | ManualFailure | LinkFailure | ActivationFailure	Identifier Mode IfList IfDownList IfCopyList
Sensor Status Change: Power Supply	OK=1 NoPower = 6 (NotPowered or PowerFailure) NotInstalled = 3	SensorName
Sensor Status Change: Fan	OK = 1 Warning = 13 Critical = 14	SensorName SensorValue
Sensor Status Change: Temperature	OK = 1 Warning = 10 Critical = 11	SensorName SensorValue
Resource Utilization Status Change: CPU/Memory/Disk	OK = 1 High = 2	SensorName SensorValue

Notes:

- The Blue Coat Sensor MIB does not define a separate **PowerFailure** status code for the Power Supply sensor status change; the **NoPower** status code is used with the trap. The WebUI does distinguish between “not powered” and “power failure”.

- Not all appliance fan or temperature sensors report a Warning status code. See Sensor Thresholds, page 8-1 for information about sensor threshold values.
- The Resource Utilization status change reports a High status code if the **UsagePercent** value is equal to or greater than 90%.

Remote Logging

Use **Remote Logging** to send appliance **System Log** and/or **Session Log** data to remote syslog servers. This is useful in many distributed corporate environments. Edit and enable a server in the **Remote Logging** panel. Up to eight remote syslog servers can be configured. See the Overview of Common Tools, page 6-5 for information on using the tool icons.

Note Make sure the segment's **Session Log Mode** option is set to **All Sessions to Remote Syslog** or **Errors to Remote Syslog** if you want to send session log data for remote logging. The **Session Log Mode** configuration must match the **Remote Logging Options** configuration in order to send the specified logs to a remote syslog server.

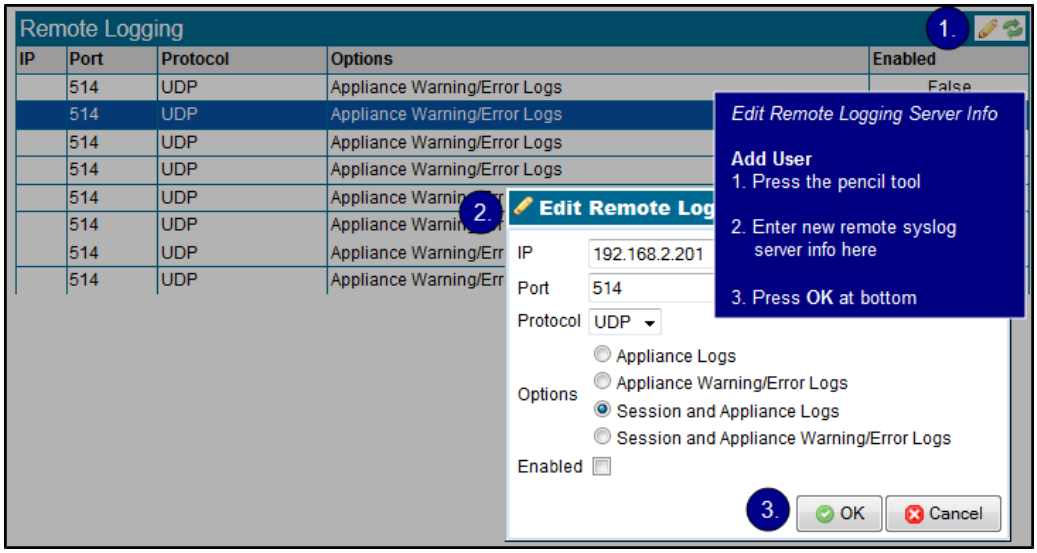

Choosing to send **Session and Appliance Logs** might result in significant traffic to the remote syslog server.

Date/Time

Use the top **Date/Time** panel to set the basic system time and date settings, as well as the time zone. In the lower panel, select whether NTP is used to synchronize the system to a network time server. See the Overview of Common Tools, page 6-5 for information on using the tool icons.

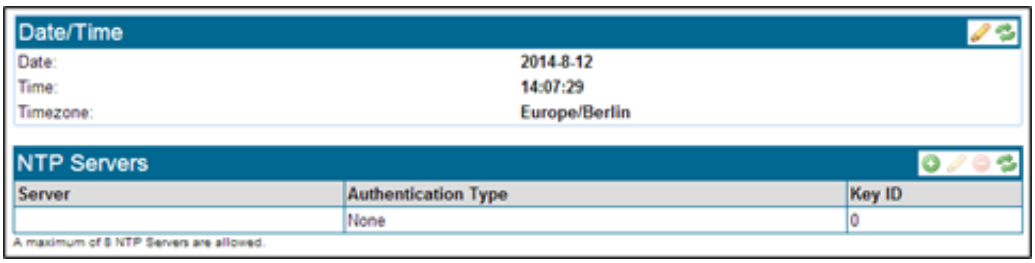

In the figure, the **Date/Time** panel shows a system that is not configured to use NTP, and which is located in the Berlin time zone. Click on the **Edit** (pencil) tool to open up an edit window where you can change the settings.

Note The system requires a reboot after changes are made to the date and time of day settings.

Configure System Date/Time and Timezone

To configure the system date and time, use the **Date/Time** option on the **(Platform Management)** menu (under the system name).

Enter Standard Date and Time

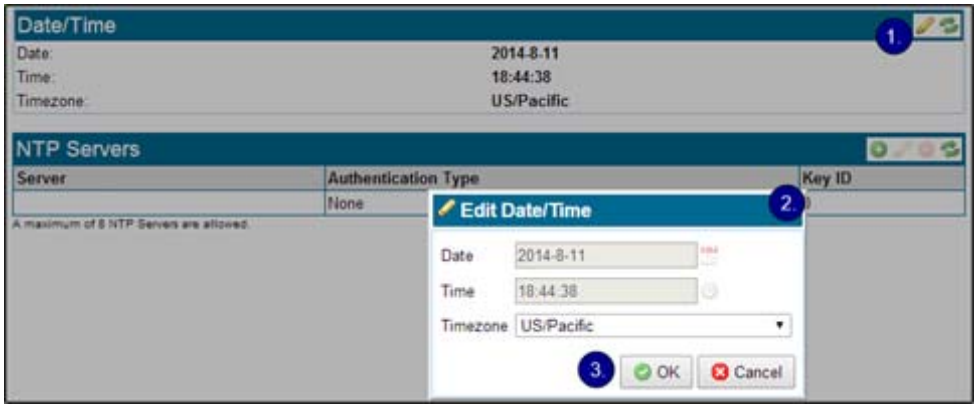

Step 1 Click **Edit**. The **Edit Date/Time** window displays.

Step 2 Enter the required **Date**, **Time**, and **Timezone**.

Step 3 Click **OK**.

If NTP is enabled, the **Date** and **Time** fields will be disabled, as these values are being set by the Network Time Protocol (NTP). In order for NTP to operate you need to configure one NTP server, and ideally, a second NTP server. NTP will not be able to resolve NTP server hostnames if there are no nameservers configured (DHCP or manually). Once the settings are configured and **OK** is clicked to save the settings, the screen will refresh.

Use Authenticated NTP Servers

Authenticated NTP synchronization can occur with an NTP server using shared key authentication. You must have the Manage Appliance user role to modify the list of NTP servers. Previously added NTP server configurations may be edited; use the **Edit** tool to access the **Edit NTP Servers** window. Sensitive newly entered and in-progress edits will be masked out for security.

The authenticated NTP server configurations will be backed up and available to restore through the **Backup/Restore** feature (seeBackup/Restore, page 6-70).

Note It may take up to several minutes for a newly added NTP server to become active.

Step 1 Click Add on the **NTP Servers** panel. The **Add NTP Server** window displays.

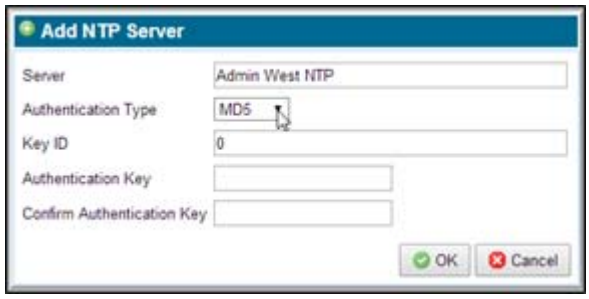

Step 2 Enter the **Server IP address** or **hostname**, select the **Authentication Type** (**None**, **MD5**, or **SHA-1**) and, if authentication is used, enter the **Authentication Key** and **Key ID**.

Note At the **Authentication Type** item, select **None, MD5**, or **SHA-1** as used by the remote NTP server (the default is **None).** An **Authentication Key** and **Key ID** are required when you select an authentication method. The entries must match the configuration on the remote NTP server.

Step 3 Click **OK**; the **Add NTP Servers** window closes.

Step 4 Click **Apply**.

Note If you have changed the date, time, NTP, or timezone, you must click **Apply** at the "Platform Config Changes" message at the bottom of the screen.

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. See the Overview of Common Tools, page 6-5 for information on using the tool icons.

The figure shows the **TACACS Servers** panel with an entry; initially the table will be empty. Click **Add** to create an entry.

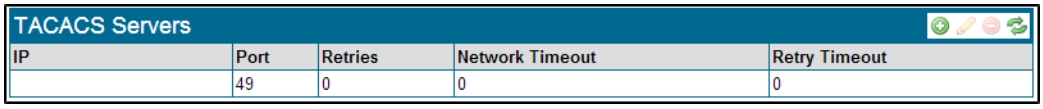

Enter the required information. The **Secret** value must match the secret value configured on the ACS server.

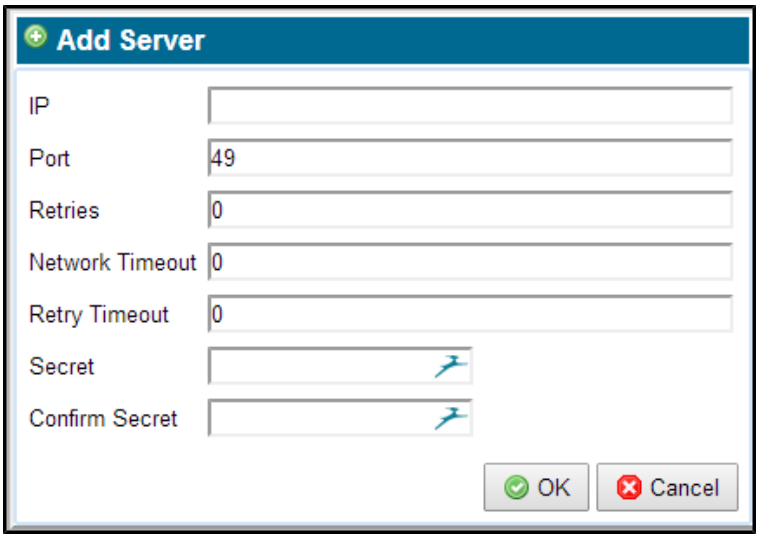

If TACACS is in use, the login pop up on the WebUI includes a drop down menu where you select **Remote** or **Local** authentication.

TACACS Administrator Privilege Mapping

The Cisco ACS lets a privilege level be stored as part of a user's profile. When the user is authenticated, the privilege level of the profile is communicated across TACACS to the SSL Appliance. As the appliance does not use privilege levels to control what an authenticated user can do, the privilege level is mapped to the roles supported by the SSL Appliance, as laid out in the next table.

TACACS Level	SSL Appliance Role
0	auditor
1	auditor + manage-appliance
2	auditor + manage-policy
3	auditor + manage-appliance + manage-policy
4	auditor + pki
5	auditor + manage-appliance + manage-pki
6	auditor + manage-policy + manage-pki
7	auditor + manage-appliance + manage-policy _ manage-pki
>8	cycle back to Level 0 and ascend again; so 8=0, 9=1, and so on

Users

The **Users** menu has a single panel with tool icons; see the <LI Link text>"Overview of Common Tools" for information on using the tool icons.

Only users with Manage Appliance or Manage PKI roles can make changes to the user accounts on the system.

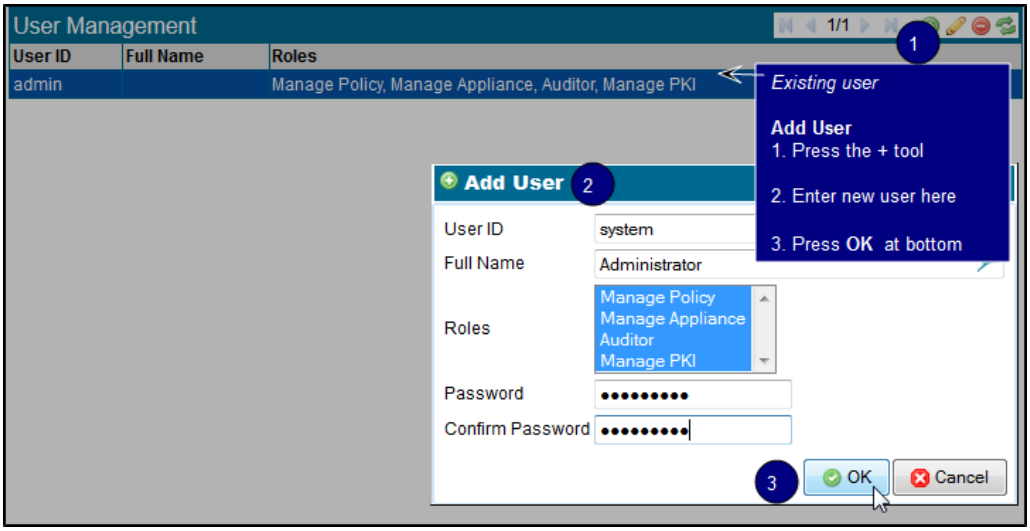

In this example, this is the **User Management** panel for a system that has one user accounts configured. More details on creating user accounts and on the meaning of different roles can be found in Configure Management Users, page 4-5.

The following table shows which actions can be performed by users with different roles:

Auditor	Manage Appliance	Manage Policy	Manage PKI	Action
			Y	Unlock secure store
Y	Y	Y	Y	View dashboards
Y	Y			View system log data
		Y	Y	View/export SSL session log, SSL errors
		Y	Y	View SSL statistics
		Y	Y	View/export intercepted certificates
			Y	Export diagnostic information: PKI state
		Y		Export diagnostic information: policy state
Y	Y	Y	Y	Export diagnostic information: platform state
		Y	Y	Export diagnostic information: SSL statistics
	Y	Y		Export diagnostic information: host statistics, NFP statistics
Y	Y	Y	Y	Export diagnostic information: platform interfaces and platform status statistics
		Y	Y	View debug information: SSL statistics
Y	Y	Y	Y	View debug information: NFE network statistics
	Y	Y		View debug information: NSM host statistics, NSM NFP statistics
		Y		Create/edit/delete rulesets, rules, segments, and user defined lists
		Y	Y	View rulesets, rules, segments, and user defined lists
		Y		Activate/deactivate segments
			Y	Create/delete/export/import internal CA keys and certificates used for resigning
			Y	Delete/import external CA certificates
			Y	Delete/import CRLs

Auditor	Manage Appliance	Manage Policy	Manage PKI	Action
			Y	Import/delete trusted certificates
			Y	Import/delete known keys and certificates
		Y	Y	View PKI information
Y	Y	Y	Y	View software, hardware details
	Y			Configure appliance settings: management network, system time, alerts
Y	Y	Y	Y	View appliance settings
	Y			Create/edit/delete user accounts
			Y	Assign/remove Manage PKI role
Y	Y		Y	View user accounts
Y	Y			View appliance settings: alerts
		Y		Backup/restore policy
			Y	Backup/restore PKI information
	Y			Backup/restore user accounts
	Y			Backup/restore platform and alert settings
	Y			Halt/reboot appliance
			Y	Import user interface certificate and key
	Y			Configure ACL by IP Address
	Y			Configure SNMPv3
		Y		Configure Host Categorization
	Y			Configure NTP Server
			Y	Configure HSM
	Y			Update the BIOS, Firmware
	Y			Configure license
Y	Y	Y	Y	Clear screen in CLI
Y	Y	Y	Y	Edit grid size in WebUI

Alerts

Use the **Alerts** panels to configure the e-mail details that the system will use to send out alerts, the events to be monitored, and the conditions under which an alert is generated.

Use the upper **Alert Mail Configuration** panel to configure details of the e-mail system. Click **Edit** to bring up the **Edit Alert Mail Configuration** window. See the Overview of Common Tools, page 6-5for information on using the tool icons.

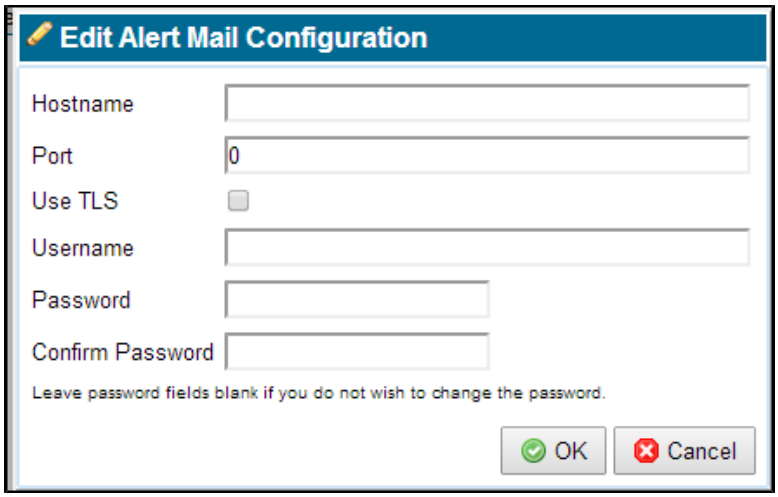

Enter the data as required:

- **Hostname**: Name or IP address of the SMTP server used to send e-mail
- **Port:** Port number on the SMTP server used to send e-mail
- **Use TLS**: Enable/disable the use of encryption (TLS) when sending e-mail
- **Username**: Username of the account used to send e-mail
- **Password**: Password for the account used to send e-mail.

Note If your enterprise is using Google Apps for e-mail, the correct SMTP Server Address is ‘aspmx.l.google.com’, not ‘smtp.gmail.com’. Ensure that DNS resolution is properly configured. Alerts can only be sent to users on the same domain with this SMTP configuration

Configure alerts on the lower **Alert Rules** panel. Each alert can be triggered by a specific set of conditions, and can be sent to one or more e-mail recipients. Click **Add** to open the **Add Rule** window and configure the rule.

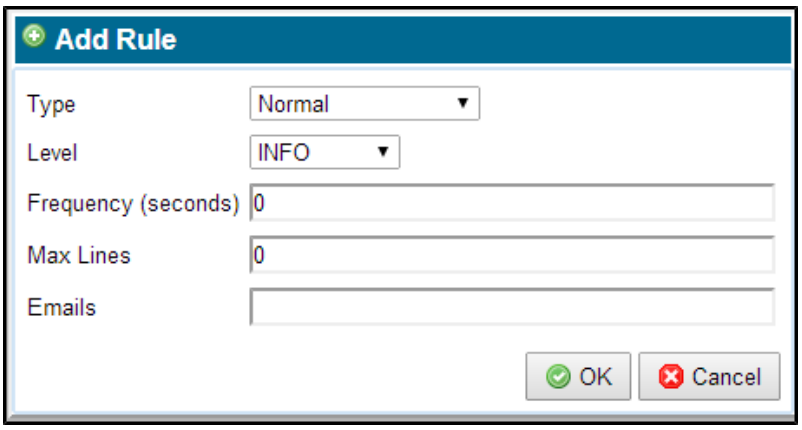

Type

- **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

Level: These levels correspond to levels associated with entries in the system log files. So, 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.

- ERROR
- FATAL
- INFO
- WARNING

Frequency (seconds): Control how frequently the alert message should be sent.

Max Lines: Control how many lines from the system log are included in the e-mail.

Emails: Specify one or more e-mail addresses; these are the users to whom the alert e-mails will be sent.

License

View and update the Host Categorization license(es).

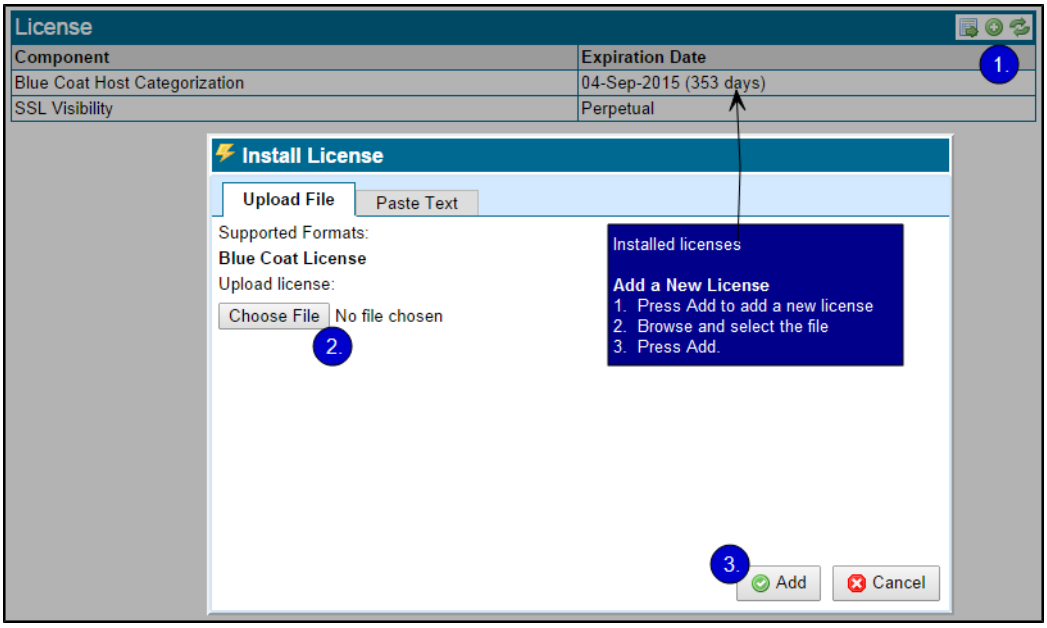

See System Status, page 4-6 for extended information on using the **License** panel. See the Overview of Common Tools, page 6-5 for information on using the tools.

Any current, active licenses appear in the **License** panel. The **License** information in the window footer will indicate the license status (depending on the state; see System Status, page 4-6 for details).

Licensing details are available in the **System Log** (View System Log Entries, page 6-9):

- If a valid license is present and not expiring within 90 days, no system log message displays
- If a valid license is present but expiring within 30 to 90 days, an INFO message displays
- If a valid license is expiring within 30 days, a WARNING message displays
- If no valid license is present, or the existing license has expired, an ERROR message displays.

License status can also be viewed on the physical LCD screen, and on the footer of the Dashboard.

Tip Configure an e-mail alert (Alerts, page 6-67) to remind yourself about a pending license expiration.

Backup/Restore

Determine the various elements of the system configuration to be saved to or restored from a remote storage system.

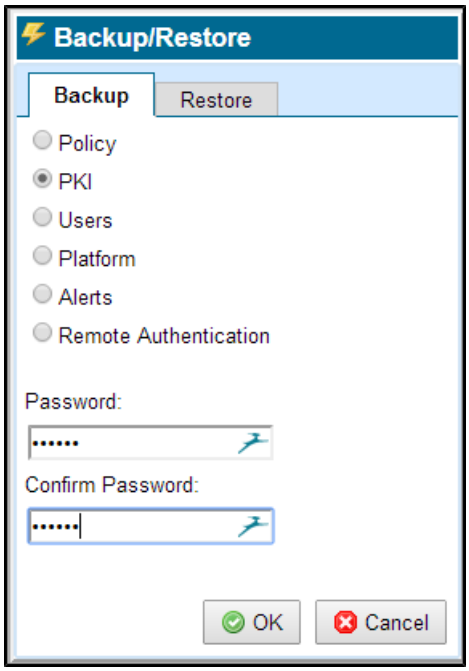
Backup/Restore
Backup
Restore
Policy
PKI
Users
Platform
Alerts
Remote Authentication
Password:
Confirm Password:
OK
Cancel

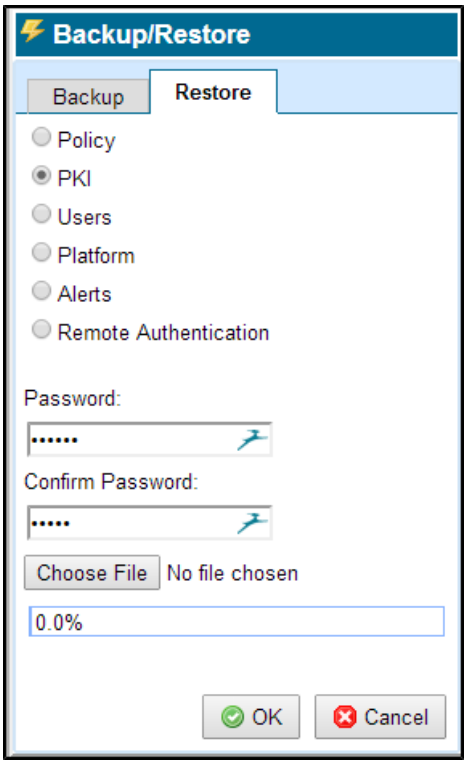
Backup/Restore
Backup
Restore
Policy
PKI
Users
Platform
Alerts
Remote Authentication
Password:
Confirm Password:
Choose File
No file chosen
0.0%
OK
Cancel

The item to be backed up or restored is indicated by selecting the radio button associated with that item. A password must be provided when backing up data and when restoring the data. All backup files created on the SSL Appliance are encrypted for security.

Note To back up or restore all configuration items, or to schedule the job, use Blue Coat Management Center 1.4.2.1 or later. For detailed information, refer to the Blue Coat Management Center Configuration Guide, which is available on BlueTouch Online Documentation (https://bto.bluecoat.com/documentation/) as a WebGuide.

Halt/Reboot

Halt or reboot the system.

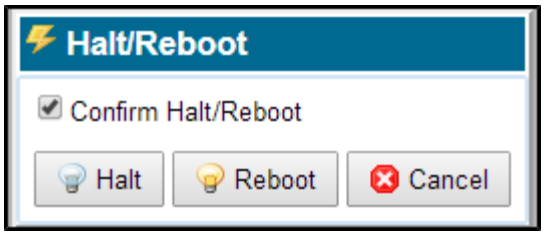

The **Confirm Halt/Reboot** check box must be checked. The Halt and Reboot buttons are grayed out until this is done.

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

Import UI Certificate/Key

Use this menu item to import a signed SSL server certificate, for use by the web server providing the WebUI management for the system. Upload the file or paste the certificate text in.

By default the system uses a self-signed server certificate which will cause warnings from browsers. See Configure the Browser, page 6-1 for details.

.

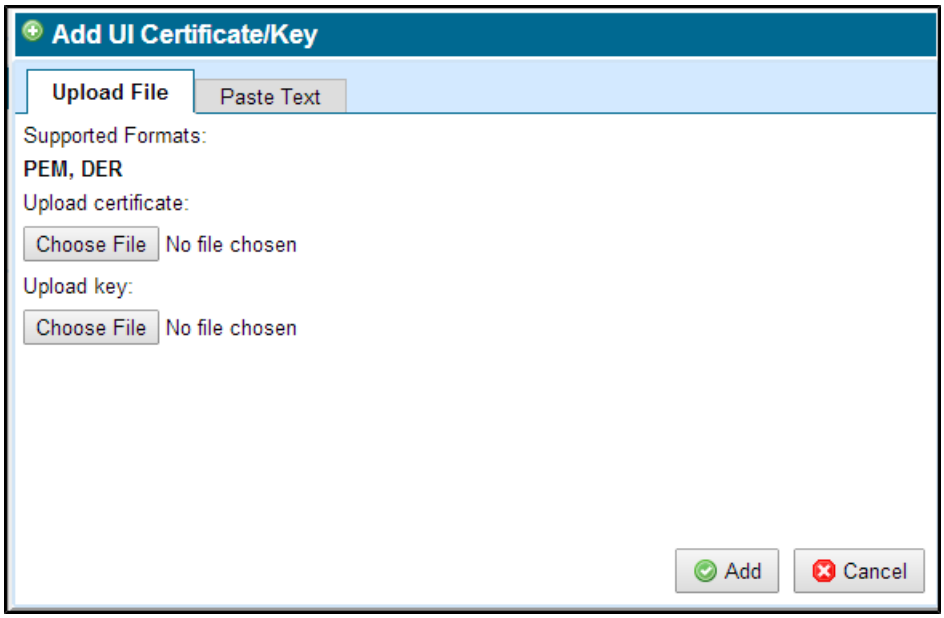

Update

Use the **Update** menu item 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. An invalid update file will not be applied.

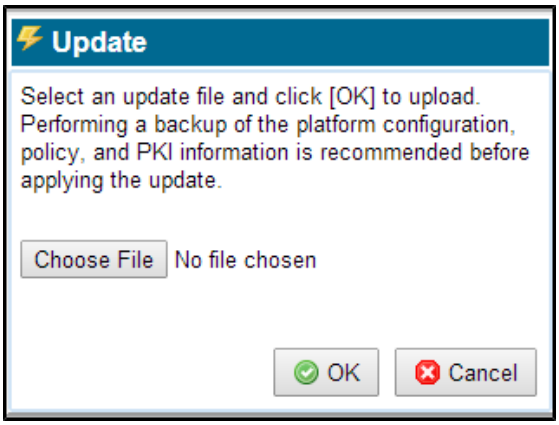

Click **Choose File** to open a window where you browse the system and select the update file to use. Click **OK**, and the file is checked, and if valid, is copied to the system and applied.

Note Once you have upgraded to software version 3.7.x or higher, the SSL Appliance cannot be downgraded without the assistance of Cisco Technical Support.

The following items will automatically be restored after the upgrade:

- System log
- Management configuration (NTP, hostname, IP, timezone, date, remote syslog)
- Policy (rulesets and segment definitions)
- PKI store
- User database
- Alerts configuration
- Remote authentication configuration

Login Banner

You may configure a statement to be presented to users during the login process, requiring the user to acknowledge and agree to the message prior to logging in.

When a Login Banner is configured, it will appear *before* the user logs on, whether using the WebUI or accessing the SSL Appliance from the Command Line Diagnostic Interface. The user is required to acknowledge the banner terms before proceeding.

Configure a Login Banner

Step 1 Start at **(Platform Management) > Login Banner.**

Step 2 Click Edit in the **Login Banner** header.

Step 3 Enter the banner content (maximum of 2047 characters) in the field on the **Login Banner** window. A configuration example is shown next:

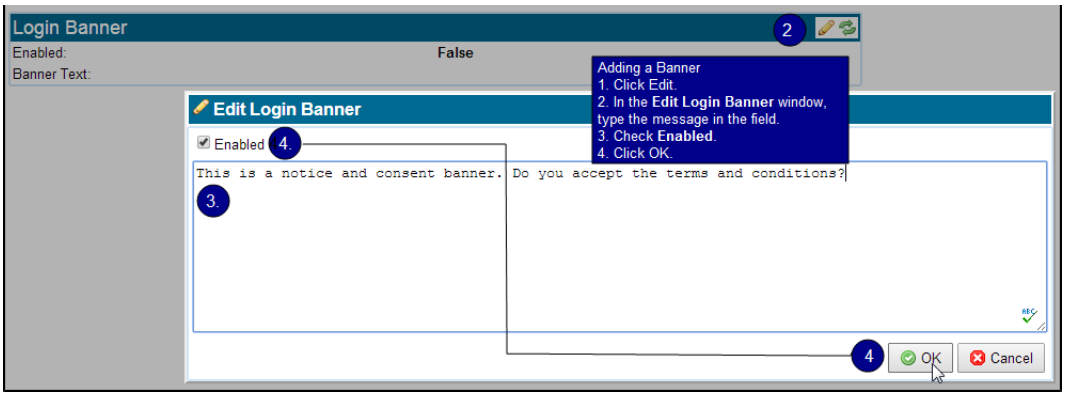

- The banner text may include a hostname variable: SV_HOSTNAME. The variable must be preceded by '${' and followed by '}', and be in capital letters. For example:

The following banner text, entered in the **Edit Login Banner** window: "You are about to login to ${SV_HOSTNAME}." will result in an actual banner display of "You are about to login to MySSLAppliance." The hostname is derived automatically. Next is an example variable configuration.

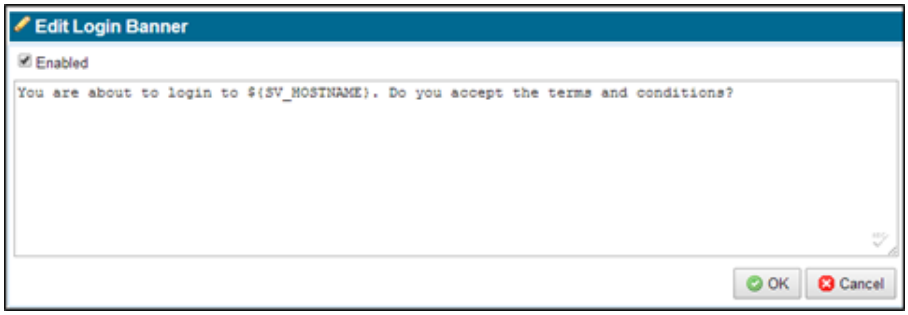

Step 4 Check **Enabled**, then click OK.

Step 5 Click **Apply** at the bottom of the screen:

Once a banner is enabled, its text will appear in the **Login Banner** window. No variable is used in this example.

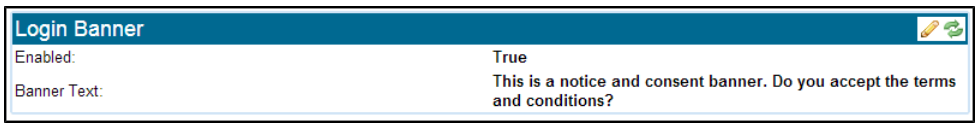

The banner is displayed when users log in to any SSL Appliance management interface. It requires all users to accept the terms and conditions before they can be authenticated and log in.

A console connection displays the banner before the password prompt.

Note It might take several seconds for a serial console to synchronize with a new banner message on start up. Hence, the first serial connection made after a change to a banner might display the wrong banner. In subsequent connections, you will see the new prompt.

On a console session, entering the password after seeing the banner affirms acceptance of the conditions.

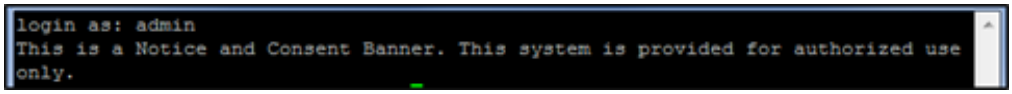

On the WebUI users must click the acceptance box.

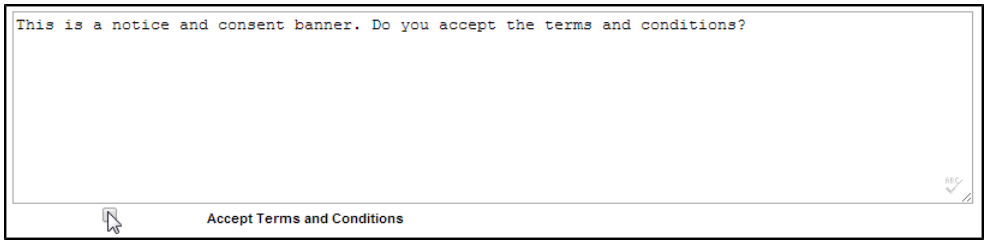

Preferences

Use the **Preferences** menu item to configure preferences that affect the WebUI screen layout. See the Overview of Common Tools, page 6-5 for information on using the tools.

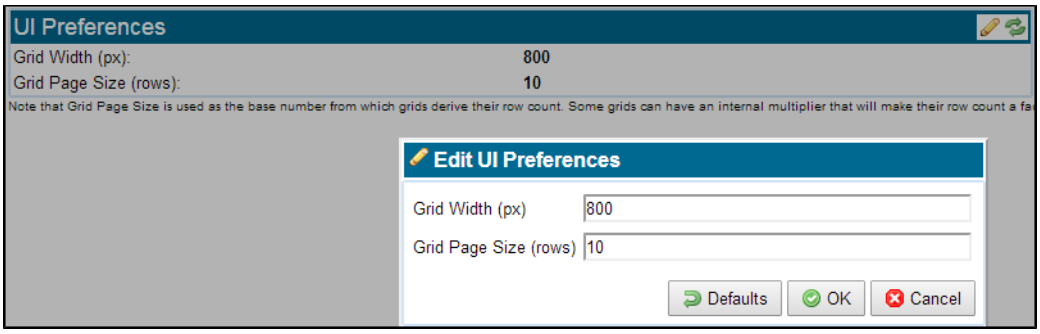

Click **Edit** to display the **Edit UI Preferences** window, also shown in the figure. Use it to change the **Grid Width and Grid Page Size** values, or to force them back to the system defaults.

Note Multistage 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 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.

User Management

Use the **User** menu, which displays at the far right of the menu bar, to change your password or log out.

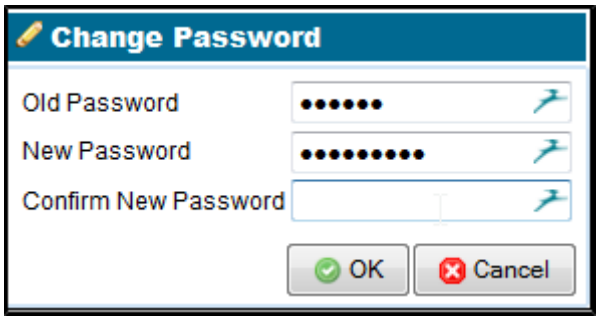

Change Password

Enter your current password, then the new password. Passwords must be a minimum of 8 characters long. It is good practice to always use at least one non-alphabetic character, at least one uppercase letter, one lowercase letter, and one digit. A user cannot reuse any of the last six passwords, including the user's current password.

Note If a Manage Appliance user edits another user's password, the action is considered a password reset and is not subject to the password reuse policy. However, if a Manage Appliance user edits their own password, the action is considered a password change, not a reset, so the newly-set password is subject to the password reuse restriction.

Logout

Selecting the **Logout** option will log the user off, and then display the login window.

Troubleshoot the System

Note Please read through all the information in this section of the document before contacting 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: Cisco ISL, MPLS, GRE, IP-in-IP, UDP, ICMP, ARP, SOCKS, DTLS, and 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. Should SSL 2.0 traffic be encountered, the SSL Appliance will either Cut Through or Reject the flow according to the Undecryptable SSL Handling parameter in the SSL Inspection Policy. SSL 2.0 ClientHello messages are supported, as long as the rest of the SSL handshake is done 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, but only if the action in the inspection policy is “Decrypt: server key is known” and RSA is used as the key exchange algorithm. The reason for this limitation is that the CertificateVerify SSL handshake message sent after the client certificate is digitally signed by a key only known to the client. The implication is that the CertificateVerify message cannot be modified, which in turn implies that 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 they use an unsupported cipher suite (Section 9.6). SSL sessions rejected because of a client certificate appear in the SSL session log with an Error event value and Reject action.

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/mask, and destination TCP port.

Supported Cipher Suites

The next table lists all the cipher suites that are supported by the SSL Appliance, and shows which can be inspected when in-line and which when in passive-tap mode. Any cipher suites that are not supported will be handled by the policies configured for undecryptable traffic.

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_Anon_WITH_RC4_128_MD5	Yes	No	0x0018
TLS_DH_Anon_WITH_DES_CBC_SHA	Yes	No	0x001A
TLS_DH_Anon_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_Anon_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_Anon_WITH_AES_256_CBC_SHA	Yes	No	0x003A
TLS_RSA_WITH_AES_128_CBC_SHA256	Yes	Yes	0x003C
TLS_RSA_WITH_AES_256_CBC_SHA256	Yes	Yes	0x003D
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA	Yes	Yes	0x0041
TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA	Yes	No	0x0045
TLS_DH_Anon_WITH_CAMELLIA_128_CBC_SHA	Yes	No	0x0046
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256	Yes	No	0x0067
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256	Yes	No	0x006B
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	Yes	Yes	0x0084
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA	Yes	No	0x0088
TLS_DH_Anon_WITH_CAMELLIA_256_CBC_SHA	Yes	No	0x0089

TLS_RSA_WITH_CAMELLIA_128_CBC_SHA256	Yes	Yes	0x00BA
TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA256	Yes	No	0x00BE
TLS_DH_anon_WITH_CAMELLIA_128_CBC_SHA256	Yes	No	0x00BF
TLS_RSA_WITH_AES_128_GCM_SHA256	Yes	Yes	0x009c
TLS_RSA_WITH_AES_256_GCM_SHA384	Yes	Yes	0x009d
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256	Yes	No	0x009e
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384	Yes	No	0x009f
TLS_DH_Anon_WITH_AES_128_GCM_	Yes	No	0x00a6
TLS_DH_Anon_WITH_AES_256_GCM_SHA384	Yes	No	0x00a7
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA256	Yes	Yes	0x00C0
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA256	Yes	No	0x00C4
TLS_DH_anon_WITH_CAMELLIA_256_CBC_SHA256	Yes	No	0x00C5
TLS_ECDHE_ECDSA_WITH_NULL_SHA	Yes	No	0xC006
TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	Yes	No	0xC007
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	Yes	No	0xC008
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	Yes	No	0xC009
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	Yes	No	0xC00A
TLS_ECDHE_ECDSA_WITH_CAMELLIA_128_CBC_SHA256	Yes	No	0xC0072
TLS_ECDHE_ECDSA_WITH_CAMELLIA_256_CBC_SHA384	Yes	No	0xC0073
TLS_ECDHE_RSA_WITH_CAMELLIA_128_CBC_SHA256	Yes	No	0xC0076
TLS_ECDHE_RSA_WITH_CAMELLIA_256_CBC_SHA384	Yes	No	0xC0077
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_Anon_WITH_NULL_SHA	Yes	No	0xC015
TLS_ECDH_Anon_WITH_RC4_128_SHA	Yes	No	0xC016
TLS_ECDH_Anon_WITH_3DES_EDE_CBC_SHA	Yes	No	0xC017
TLS_ECDH_Anon_WITH_AES_128_CBC_SHA	Yes	No	0xC018
TLS_ECDH_Anon_WITH_AES_256_CBC_SHA	Yes	No	0xC019
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	Yes	No	0xC023
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	Yes	No	0xC024
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	Yes	No	0xC027
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	Yes	No	0xC028
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	Yes	No	0xc02b
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	Yes	No	0xc02c

TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	Yes	No	0xc02f
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	Yes	No	0xc030
TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256	Yes	No	0xcc13
TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_	Yes	No	0xcc14
TLS_DHE_RSA_WITH_CHACHA20_POLY1305_SHA256	Yes	No	0xcc15
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 no support for static DH (Diffie-Hellman) key exchange, or DSS (Digital Signature Standard) authentication.

Note When operating in Passive-Tap mode there are some cipher suites that cannot be inspected: Ephemeral, Elliptic Curve and Anonymous DH key exchanges. When operating in inline modes it is possible to inspect SSL sessions using Ephemeral, Elliptic Curve 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. The next list shows supported SSL cipher suites:

Supported SSL Cipher Suites

The SSL Appliance supports inspecting SSL sessions with the following cipher suites:

- AES128-SHA256
- AES256-SHA256
- DHE-RSA-AES128-SHA256
- DHE-RSA-AES256-SHA256
- AES128-SHA
- AES256-SHA
- DHE-RSA-AES128-SHA
- DHE-RSA-AES256-SHA
- TLS_ECDHE_ECDSA_WITH_CAMELLIA_128_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_CAMELLIA_256_CBC_SHA384
- TLS_ECDHE_RSA_WITH_CAMELLIA_128_CBC_SHA256
- TLS_ECDHE_RSA_WITH_CAMELLIA_256_CBC_SHA384

There are no restrictions on cipher suites for policies with actions that do not involve inspecting the traffic. So, it is fine to have a policy that prevents SSL traffic using static DH from setting up connections across the network for example.

Support for SSL Record Layer Compression

The SSL specification allows for SSL record layer compression using 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 will provide the required information. If you can see the counts for detected SSL session increasing then 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 will provide 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 **SSL Session Log**: ensure that you are viewing the first page of session log data, then click **Refresh**. 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 **SSL Session Log** indicates which segment and entry is for, so you need to know the segment ID that is associated with the segment you are troubleshooting. Find this on the **Policies > Segment** screen.

Verify that the Inspection Policy is Set Up Correctly

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 Configure Rulesets to Handle SSL Traffic, page 6-18 for more details.

Known Server vs 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.

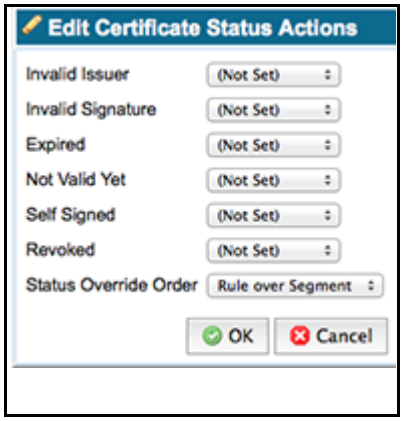

Do not install server certificates in the Trusted Certificates store if you have the private key for that server: those certificates belong in the Known Certificates and Keys store.

The Trusted Certificates store is only used to solve specific certificate validation problems, that is,. trusting self-signed certificates or trusting certificates for which you don't want to install the CA certificate chain. Refer to PKI Management, page 6-42.

Caveats when Enabling/Disabling SSL Inspection

Immediately after you connect a segment to the network or activate inspection, it might 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 Domain Name, 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.

A 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 e-mail clients, cache the cryptographic state and might re-use 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 re-using the same session. If the full handshake occurred before the appliance was put inline, it cannot decrypt flows re-using that session. Most servers allow sessions to be re-used only for a few hours, after which they force clients to establish new sessions. Therefore, the **SSL Session Log** might

show Uncached sessions for a few hours after installing the device on the network or activating inspection. As soon as 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.

Another caveat is that SSL clients might report SSL session failures if you disconnect the SSL Appliance. If an application, for example, Microsoft Outlook, supports SSL session re-use, it will report a failure when it tries to re-use the SSL session. The reason this fails is that when the full SSL handshake was used to establish the initial SSL session the SSL appliance was inline and acting as a man in the middle (MITM). So the session that the client has saved and is trying to re-use was actually a session from the client to the SSL Appliance rather than to the server. The client does not know this as 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 as it does not recognize it.

Generating the Resigning CA Certificates

Inspecting SSL sessions in any of the inline modes requires at least one resigning CA certificate and private key, unless only Known Key decryption is used. The SSL Appliance can generate the resigning CA resigning certificate authorities 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 Install a Local CA for Certificate Resign, page 4-8 and Resigning Certificate Authorities, page 6-43 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 might be displayed by Internet Explorer and the update service will fail.

This error occurs because the CA of the certificate presented by the update website server, is found not to be a Microsoft server, and thus 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. Windows update services should now function normally.

Note A default list of certificate Common Names (CNs) for sites that it is not possible to inspect traffic to are included in the DN list menu. A rule using this list can be added to a ruleset to ensure that traffic to these sites is not inspected.

Issues with Alerts

- If you fail to receive e-mail alerts, check the system log file for errors. The following might also prevent e-mail from being sent or delivered:
- If your SMTP server requires authentication, check that 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 e-mail addresses are correct.
- If your enterprise is using Google Apps for e-mail then the correct SMTP Server Address is 'aspmx.l.google.com', not 'smtp.gmail.com'. Ensure that DNS resolution is properly configured. 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, page 6-14 for details on how to generate diagnostic files.

The support engineers might request further diagnostic information such as SSL statistics, non-SSL statistics, and the **SSL Session Log** (if enabled). The engineers will not request a copy of the PKI store because it might 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 this can commence 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 then appropriate directions will be given by the support engineer.

Command Line Diagnostics Interface

You might be asked to use the Command Line Diagnostics interface via an SSH or serial console connection, by Customer Service, to aid in troubleshooting (see the *Getting Started Guide* for details). The following table lists each command, and the related action.

In this table, ">" indicates a sub-command.

Enter `?` for a list of commands.

Enter `(command)?` For a list of related commands. For example, `platform?` returns `platform halt` and `platform reboot`.

Command	Action
bios update	Update the BIOS
capture reset	Reset the network capture state and remove all captures stored on disk
capture select	Select capture mode and interfaces; optionally filter by source or destination IP address **Parameters**: [src-ip] <IP_address> - captures only traffic from this source IP on the specified interface [dst-ip] <IP_address> - captures only traffic to this destination IP on the specified interface <interface_#> - Interface to capture on (required) <interface_#2> ... <interface_#32> - Additional interfaces to capture on (optional)
capture start	Start capturing network traffic

Command	Action
capture status	Show the current network capture status **Example**: admin> capture status Ready to configure and start capture Selected interface(s) : 5 Selected mode : normal
capture stop	Stop capturing network traffic
challenge show	Show backend authentication challenge
clear	Clear screen
counters interface	Show external interface port statistics counters **Parameters**: [diff] <seconds> - Show difference in counters over specifed time in seconds (optional) <interface_#> - Interface to show port statistics for (required) <interface_#2> ... <interface_#4> - Additional interfaces to show counters for (optional) **Flags (optional):** [nonzero] - Show only non-zero counters
counters npu	Show NPU counters **Parameters**: [diff] <seconds> - Show difference in counters over specifed time in seconds (optional) card <card_#> - NPU card to show counters for; default = 1, if not specified **Flags (optional):** [nonzero] - Show only non-zero counters
counters packets	Show packet counters **Parameters**: [diff] <seconds> - Show difference in counters over specifed time in seconds (optional) **Flags (optional):** [nonzero] - Show only non-zero counters
counters ssl	Show SSL counters **Parameters:** [diff] <seconds> - Show difference in counters over specifed time in seconds (optional) **Flags (optional):** [nonzero] - Show only non-zero counters

Command	Action
counters switch	Show switch counters (port statistics on the NFE interface) **Parameters**: [diff] <seconds> - Show difference in counters over specifed time in seconds (optional) [card] <card_#> - NPU card to show switch counters for; default = 1, if not specified **Flags (optional):** [nonzero] - Show only non-zero counters
counters tcp	Show TCP counters **Parameters:** [diff] <seconds> - Show difference in counters over specifed time in seconds (optional) **Flags (optional):** [nonzero] - Show only non-zero counters
debug crash show	Display SCP command string to retrieve last crash dump
debug disk	Display disk usage statistics
debug dmesg	Display kernel message (dmesg) output
debug eth0	Display management network interface information
debug firewall	Display IPv4 firewall rules and statistics
debug firewall6	Display IPv6 firewall rules and statistics
debug memory	Display memory statistics
debug netstat	Display network connection information
debug npu queue	Display NPU queue information
debug ping	Ping an IPv4 host. **Parameters:** Host name or IPv4 address (required)
debug ping6	Ping an IPv6 host. **Parameters:** Host name or IPv6 address (required)
debug proc	Display running processes
debug relays	Display relay status
debug resolve	Look up IP address of a host **Parameters**: Host name (required)
debug route	Display routing information
debug switch	Display switch status
diags reset	Reset diagnostics state

Command	Action
diags select	Select options for diagnostics collection **Parameters**: [from] <MM-DD-YYYY> - Start date for statistics history [to] <MM-DD-YYYY> - End date for statistics history **Flags (optional):** At least one flag must be specified. [pki] - Include PKI store summary [policy] - Include policy [platform] - Include platform state [ssl-stats] - Include SSL statistics history [host-stats] - Include host statistics history [npu-stats] - Include NPU statistics history [platform-stats] - Include platform statistics history [interface-stats] - Include external interface statistics history
diags start	Start diagnostics collection
diags status	Check diagnostics status
error	Translate error codes **Parameters**: <error code> (required)
error counts	Dump flow error codes and counts
exit	Logout [quit] is an alias for the exit command.
fips mode show	Show FIPS/non-FIPS security mode
license add	Install a new license, overwriting any currently installed license Requires the license to be pasted during execution. **Example:** admin> license add Paste SSLV license and enter Ctrl+D when done: <license string> Ctrl+d
license export	Export the currently installed license.
license status	Show license status
master key create	Create master key, phase II (available only during bootstrap phase)
master key pin	Use PIN protection (available only during bootstrap phase)
master key settings	Create master key, phase I (available only during bootstrap phase)
master key storage	Set physical location of master key (local or USB)

Command	Action
network set ip	Set management network static IPv4 configuration **Parameters:** <IPv4 address/netmask> (required) [netmask] - Netmask (default:255.255.255.0) [gateway] <IP address> - Default gateway
network set ip dhcp	Enable DHCP in management network IPv4 configuration
network set ip disabled	Disable IPv4 in management networks
network set ip6	Set management network static IPv6 configuration **Parameters:** <IPv6 address/netmask> (required) [netmask] - Netmask (default:ffff:ffff:ffff:ffff::) [gateway] <IP address> - Default gateway
network set ip6 dhcp	Enable DHCP in IPv6 management network configurations
network set ip6 disabled	Disable IPv6 in management network configuration
network set ip6 slaac	Enable SLAAC in management network IPv6 configuration **Flags (optional):** [stateless-dhcp] - Use stateless DHCP for nameserver assignment
network set mtu	Set management network MTU **Parameters:** <MTU value> (required)
network show	Show network IP and MTU settings
network-acl edit	Edit the IPv4 access control list; enter this command to access the edit subcommands
network-acl edit > append	Add a new rule to the end of the pending IPv4 access control list **Parameters:** [allow] or [block] - Specify the rule action (required) <IP address/netmask> or [all] - specify the address to match on (required) **Flags (optional):** [management] - Make this rule apply to management traffic [network-utilities] - Make this rule apply to network-utility traffic (e.g. ping) [snmp] - Make this rule apply to SNMP traffic
network-acl edit > clear	Clear screen
network-acl edit > commit	Commit the pending edits to the IPv4 access control list
network-acl edit > delete	Delete a rule from the pending IPv4 access control list **Parameters:** <rule #> - Specify the rule to delete
network-acl edit > error	Translate error codes **Parameters:** <error code> (required)

Command	Action
network-acl edit > exit	Exit the IPv4 access control list edit level
network-acl edit > insert	Insert a new rule into the pending IPv4 access control list **Parameters**: <rule #> - Specify where to insert the new rule (required) [allow] or [block] - Specify the rule action (required) <IP address/netmask> or [all] - specify the address to match on (required) **Flags (optional):** [management] - Make this rule apply to management traffic [network-utilities] - Make this rule apply to network-utility traffic (e.g. ping) [snmp] - Make this rule apply to SNMP traffic
network-acl edit > show active	Show the currently running IPv4 access control list
network-acl edit > show pending	Show the pending IPv4 access control list; use commit to make pending rules active
network-acl show active	Show the currently running IPv4 access control list
network-acl6 edit	Edit the IPv6 access control list; enter this command to access the edit subcommands
network-acl6 edit > append	Add a new rule to the end of the pending IPv6 access control list **Parameters:** [allow] or [block] - Specify the rule action (required) <IP address/netmask> or [all] - specify the address to match on (required) **Flags (optional):** [management] - Make this rule apply to management traffic [network-utilities] - Make this rule apply to network-utility traffic (e.g. ping) [snmp] - Make this rule apply to SNMP traffic
network-acl6 edit > clear	Clear screen
network-acl6 edit > commit	Commit the pending edits to the IPv6 access control list
network-acl6 edit > delete	Delete a rule from the pending IPv6 access control list **Parameters:** <rule #> - Specify the rule to delete
network-acl6 edit > error	Translate error codes **Parameters:** <error code> (required)
network-acl6 edit > exit	Exit the IPv6 access control list edit level

Command	Action
network-acl6 edit > insert	Insert a new rule into the pending IPv46access control list **Parameters**: <rule #> - Specify where to insert the new rule (required) [allow] or [block] - Specify the rule action (required) <IP address/netmask> or [all] - specify the address to match on (required) **Flags (optional):** [management] - Make this rule apply to management traffic [network-utilities] - Make this rule apply to network-utility traffic (e.g. ping) [snmp] - Make this rule apply to SNMP traffic
network-acl6 edit > show active	Show the currently running IPv6 access control list
network-acl6 edit > show pending	Show the pending IPv6 access control list; use commit to make pending rules active
network-acl6 show active	Show the currently running IPv6 access control list
platform factory-reset schedule	(For use by Cisco personnel only) Reset the appliance to factory state on next reboot
platform factory-reset cancel	(For use by Cisco personnel only) Cancel a scheduled factory reset
platform halt	Halt the appliance
platform packages	(For use by Cisco personnel only) Display package information
platform reboot	Reboot the appliance
platform show model	Display hardware model
platform show ntp	Show the NTP status
platform signing key	(For use by Cisco personnel only) Show name of key with which the image is signed
segment	Show details about an activated segment **Parameters:** <segment ID> (A, B, C) - Segment to show (required) <segment ID> - Second segment to show <segment ID> - Third segment to show <segment ID> - Fourth segment to show
segment all	Show details about all activated segments
segment fail all	Fail to wire the interfaces of all activated segments
segment fail	Fail to wire the interfaces of an activated segment **Parameters:** <segment ID> (A, B, C) - Segment to fail (required) <segment ID> - Second segment to fail <segment ID> - Third segment to fail <segment ID> - Fourth segment to fail

Command	Action
segment interfaces	Show statistics for all external interfaces assigned to an activated segment **Parameters:** <segment ID> (A, B, C) - Segment to show external interface statistics for **Flags (optional):** [nonzero] - Show non-zero statistics counters only
segment list	Show the status of all activated segments
segment unfail	Unfail the interfaces of an activated segment **Parameters:** <segment ID> (A, B, C) - Segment to unfail <segment ID> - Second segment to unfail <segment ID> - Third segment to unfail <segment ID> - Fourth segment to unfail
segment unfail all	Unfail the interfaces of all activated segments
session log export	Export SSL session logs; the exported file is in .tgz format **Parameters:** [start] <YYYY-MM-DD> or <YYYY-MMDD, hh:mm:ss> - Set start date (required) and time (optional) [end] <YYYY-MM-DD> or <YYYY-MMDD, hh:mm:ss> - Set end date (required) and time (optional)
session log export reset	Reset the session log export state, and delete the archived session log
snmp show	Show SNMP configuration parameters
update reset	Reset the update state and cancel any pending updates
update status	Show the current update status
uptime	Display system uptime; length of time since the appliance was last restarted or reset.
user add	Add a user **Parameters:** <user ID> - User ID to add (required) [name] <user name> - User full name (optional) [password] <password> - User password (required) **Flags (optional):** [manage-pki] - Add **Manage PKI** role [manage-appliance] - Add **Manage Appliance** role [manage-policy] - Add **Manage Policy** role [audit] - Add **Auditor** role

Command	Action
user add role	Add a role to an existing user **Parameters:** <user ID> - User ID to add role to **Flags (optional):** [manage-pki] - Add **Manage PKI** role [manage-appliance] - Add **Manage Appliance** role [manage-policy] - Add **Manage Policy** role [audit] - Add **Auditor** role
user change password	Change an existing user's password **Parameters:** <user ID> - User ID to change password of (required) The system prompts for the old user password and new user password.
user list	List all existing users
user remove	Remove a user **Parameters**: <user ID> - User ID to remove
user remove role	Remove a role from an existing user **Parameters:** <user ID> - User ID to remove role from **Flags (optional):** [manage-pki] - Remove **Manage PKI** role [manage-appliance] - Remove **Manage Appliance** role [manage-policy] - Remove **Manage Policy** role [audit] - Remove **Auditor** role
user set name	Set a user's full name **Parameters:** [name] <user name> - User full name <user ID> - User ID to set full name of
user show	Display user information **Parameters:** <user ID> - User ID to show data for
version	Display version information and serial number

Additional Screens in the WebUI

The following WebUI screens can display additional output or provide advanced options in response to an “advanced” parameter in the screen’s URL.

- Monitor Menu

  - • SSL Session Log screen
  - Debug screen
- Platform Management (system name) Menu
  - License screen

To access the advanced features of a screen, append "?advanced" to the screen's URL. For example: https://192.168.xx.xxx/#monitor.ssl_session_log?advanced

SSL Session Log

The **SSL Session Log** screen advanced view provides access to additional filters. Filtering is available for **Source IP/Mask**, **Destination IP/Mask**, **Source Port**, and **Destination Port**. You can also filter by **SSL Version**, **Flags**, and **Cipher Suite** by providing an appropriate hex value.

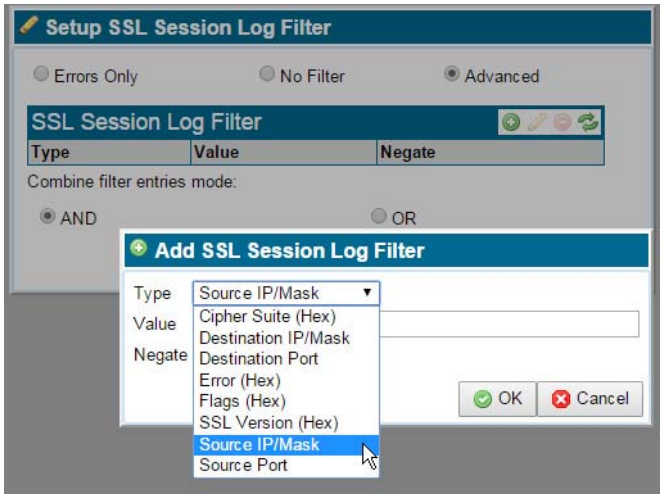

The **Negate** option excludes entries matching the specified filter. You can combine multiple filter entries using Boolean AND/OR logic. If you use multiple filter terms in an OR operation, any entry in the combination constitutes a match. To filter on multiple terms, the AND operator requires a match on all entries.

SSL Version

The SSL Version filter accepts the following hex values:

Value	Version
0300	SSL3
0301	TLS 1.0
0302	TLS 1.1
0303	TLS 1.2

Cipher Suite

Hex values for Cipher Suites can be obtained from the relevant RFCs or the TLS Cipher Suite Registry maintained by the Internet Assigned Numbers Authority (IANA). Refer to http://www.iana.org/assignments/tls-parameters/tlsparameters.xhtml#tls-parameters-4.

For example, TLS_SRP_SHA_WITH_AES_256_CBC_SHA corresponds to the values 0xC0, 0x20; the hex entry is C020.

Flags

The following flags are supported for filtering the SSL Session Log. Enter the hex value to filter by the flag.

Value	Flag
0000000000000010	CACHED_WITH_SESSID
0000000000000020	CACHED_WITH_SESSTKT
0000000000000400	FULL_HANDSHAKE
0000000000000800	REUSED_SESSION
0000000080000000	CLIENT_CERT_NOT_SUPPORTED
2000000000000000	ERROR_IN_PROCESSING

Additional filter values (for Errors, for example) may be provided by Cisco support personnel.

Debug

The Debug information that can be displayed in the WebUI is primarily intended to aid troubleshooting. The Debug panels enable users to provide information to Cisco support personnel when requested.

The default Debug screen displays the NFE Network Statistics panel. SeeDebug, page 6-15 for information about the default display.

The advanced Debug view adds three additional panels.

NSM Host Statistics

This panel shows counters for the NSM software that is running on the x86 processing complex.

NSM NFP Statistics

This panel shows counters for the NSM software that is running on the NFP chip(s).

SSL Statistics

This panel shows counters for SSL traffic processing events.

License

The License screen advanced view includes a **Delete** button that deletes the current license. Use this function only if instructed to do so by Cisco support personnel.

Sensor Thresholds

The following tables show the sensor threshold values that are monitored as alarm thresholds for SSL Appliances, for each appliance model.

- The system generates a Critical (red) alarm when a threshold value reaches or exceeds an Upper Critical or Lower Critical limit.
- The system generates a Caution (yellow) when a threshold value reaches an Upper or Lower Non-critical limit, but has not reached an Upper or Lower Critical limit.

Units of Measurement

The unit of measurement for fan speed sensors is RPM.

The unit of measurement for temperature sensors is degrees Celsius.

Thermal margin sensors are reported as negative values which, when increased to 0, will cause the CPU to throttle down or halt.

Table 8-1 SSL 1500 Thresholds

Sensor	Lower Critical	Lower Non-Critical	Upper Non-Critical	Upper Critical
Fan 1A	2000	6000	18000	20000
Fan 2A	2000	6000	18000	20000
Fan 3A	2000			
Fan 4A	2000			
Fan 4B	2000			
Power Supply Fan	0 (Stopped)			
VRD Temp				90
PCH Temp				90
Inlet Temp				40
CHA DIMM 0 Temp				87
CHA DIMM 1 Temp				87
CHA DIMM 2 Temp				87

Table 8-1 SSL 1500 Thresholds

Sensor	Lower Critical	Lower Non-Critical	Upper Non-Critical	Upper Critical
CHB DIMM 0 Temp				87
CHB DIMM 1 Temp				87
CHB DIMM 2 Temp				87
NFP Temp				90
Power Supply Temp				80
CPU Thermal Margin				0

Note Thermal margin sensors are reported as negative values which, when increased to 0, will cause the CPU to throttle down or halt.

Table 8-2 SSL 2000 Thresholds

Sensor	Lower Critical	Lower Non-Critical	Upper Non-Critical	Upper Critical
Fan Mod 1 Inlet	975	1950		
Fan Mod 1 Outlet	536	938		
Fan Mod 2 Inlet	975	1950		
Fan Mod 2 Outlet	536	938		
Fan Mod 3-5	536	938		
Power Supply Fan	0 (Stopped)			
Baseboard Temp	5	10	61	66
Front Panel Temp	0	5	36	40
NFP Temp				90
Power Supply Temp				80
IOH Thermal Margin				0
Mem P1 Thermal Margin				0
P1 Thermal Margin				0
P2 Thermal Margin				0

Note Thermal margin sensors are reported as negative values which, when increased to 0, will cause the CPU to throttle down or halt.

Table 8-3 SSL 8200 Thresholds

Sensor	Lower Critical	Lower Non-Critical	Upper Non-Critical	Upper Critical
Processor 1 Fan	968	1936		
Processor 2Fan	968	1936		
Memory Fan 1	980	1470		
Memory Fan 2	980	1470		
System Fan 1	968	1936		
System Fan 2	980	1470		
Power Supply Fan	0 (Stopped)			
Baseboard Temp	5	10	61	66
Front Panel Temp	0	5	36	40
NFP Temp				90
Power Supply Temp				80
IOH Thermal Margin				0
Mem P1 Thermal Margin				0
Mem P2 Thermal Margin				0
P1 Thermal Margin				0
P2 Thermal Margin				0

Note Thermal margin sensors are reported as negative values which, when increased to 0, will cause the CPU to throttle down or halt.

Safety Information

In addition to the information below you should read the separate Safety Notice included in the SSL Appliance packaging.

Safety Instructions

Please read all of the following instructions regarding the Cisco SSL Appliance carefully.

- Ventilation

  The Cisco SSL Appliance vents (on the front panel) and the fan openings 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

Caution 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 a dual redundant power supply that is powered by two separate power cords. Always disconnect BOTH cords to remove power from the unit.

Warning **To reduce the risk of electrical shock, do not disassemble this product. Return it to Cisco when service or repair work is required. Opening or removing covers might expose the user to dangerous voltage or other risks. Incorrect assembly can cause electric shock when this appliance is subsequently used.**

Note Opening the cover will void the warranty!

Technical Support

Thank you for choosing the Cisco SSL Appliance.

Cisco Support

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information about Cisco SSL Appliances, see **What's New in Cisco Product Documentation** at: http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html.

Subscribe to **What's New in Cisco Product Documentation**, which lists all new and revised Cisco technical documentation, as an RSS feed and deliver content directly to your desktop using a reader application. The RSS feeds are a free service.

If you have any questions or require assistance with the Cisco SSL Appliance, you can also contact Cisco Support:

- Visit the Cisco Support site at http://support.cisco.com/.
- Email Cisco Support at tac@cisco.com.
- Call Cisco Support at 1.408.526.7209 or 1.800.553.2447.