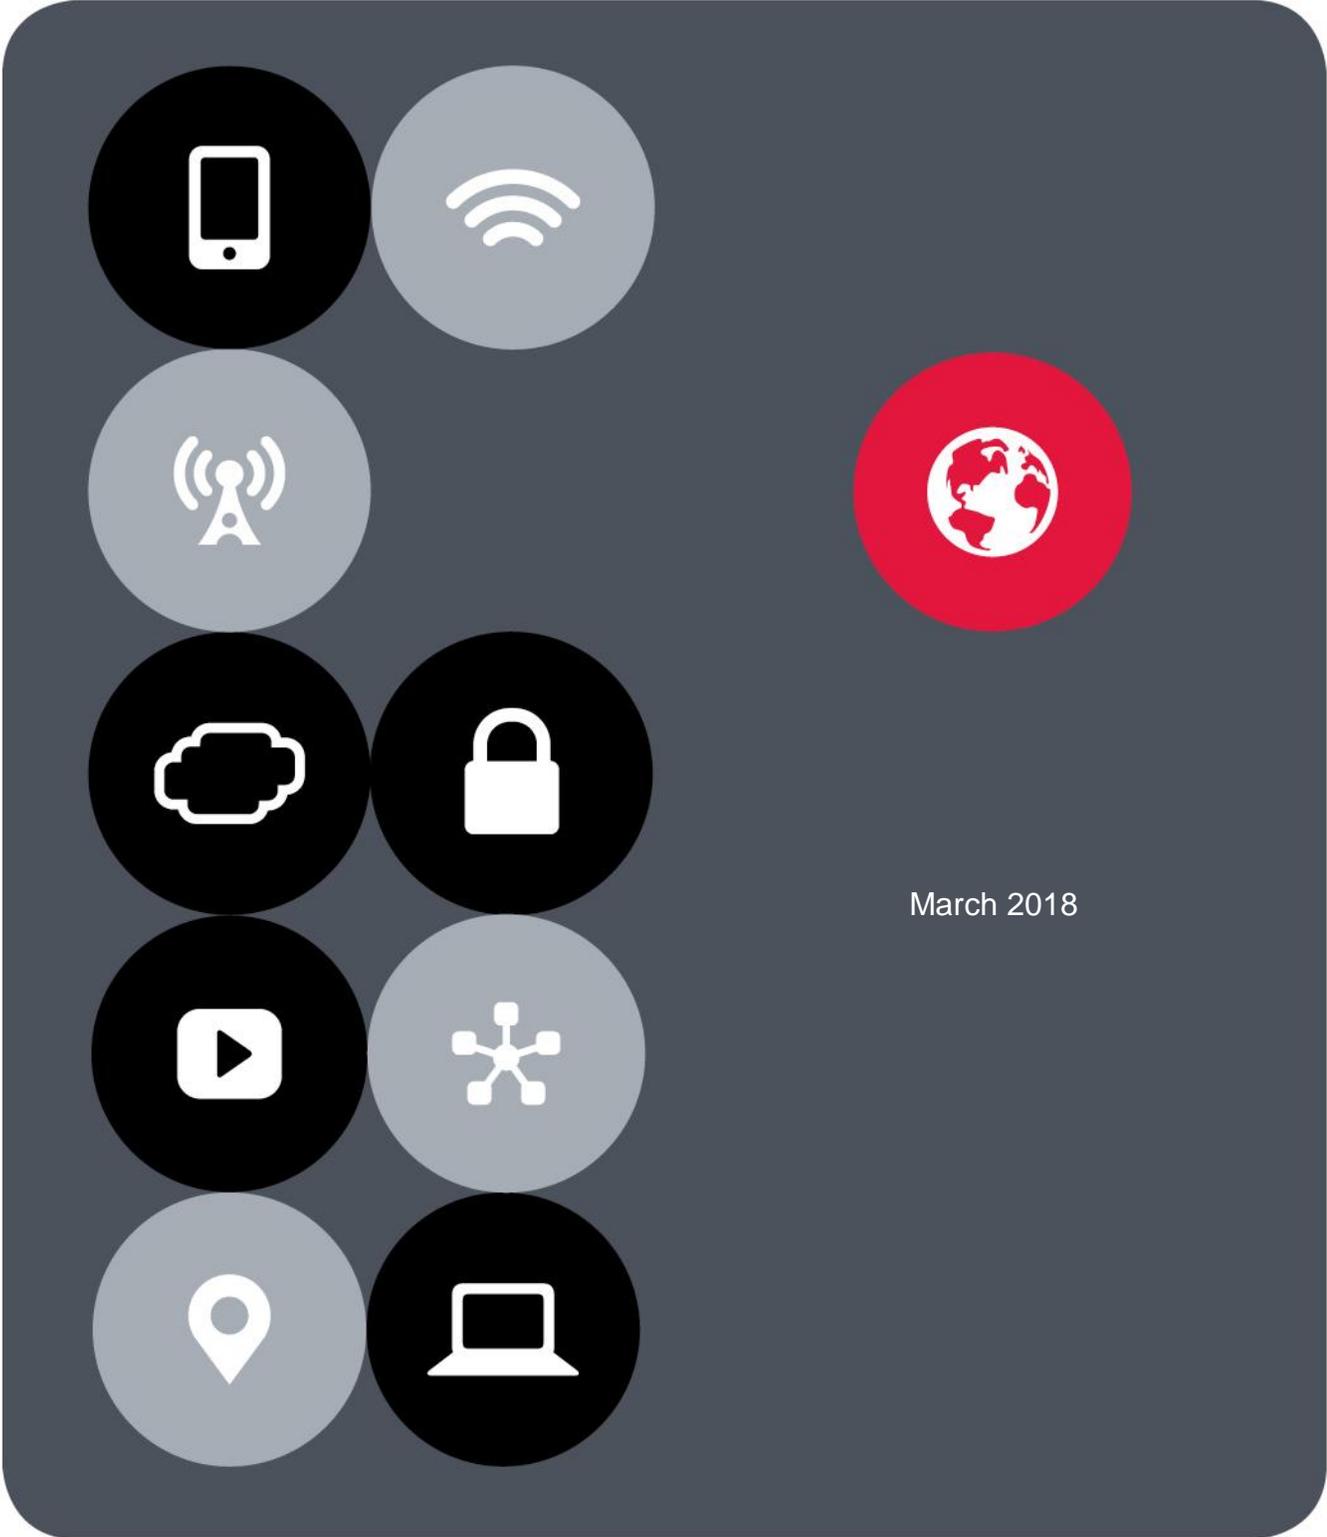




RECOMMENDED DEPLOYMENT PRACTICES

The F5 SSL Orchestrator and Cisco Firepower Solution: SSL Visibility with Service Chaining for Advanced Malware Protection



March 2018

Contents

Introduction	3
The Integrated Solution	3
SSL visibility: How do we do it?	4
SSL orchestration using security service chains	5
Deployment Planning	6
Sizing	6
License components	7
Horizontal scaling	7
Traffic exemptions for SSL inspection	8
Certificate requirements	9
IP addressing	9
Deployment modes	10
Initial Setup	11
Run the SSL Orchestrator Setup Wizard	12
Update the SSL Orchestrator version	15
Back up your F5 system configuration	17
Configuration for a Single F5 System with FirePOWER Services on Cisco ASA in L2 Mode (Burrito Design)	18
Configure SSL Orchestrator	19
Create layer 2 inline service	23
Configuration for an F5 System with FirePOWER Services on Cisco ASA in L3 Mode	25
Create the layer 3 inline service	25
Configuration for the F5 System with Cisco ASAs in TAP Mode	26
Create a receive-only service	27
Alternative Architectures	28
Two F5 systems with ASAs deployed as a service pool	28
Two F5 systems with firewalls sandwiched in the decryption zone	31
Creating service chains to link services	34
Creating TCP service chain classifier rules	34
Handling NAT	37
Testing the Solution	38

Introduction

The Secure Sockets Layer (SSL) protocol and its successor, Transport Layer Security (TLS), have been widely adopted by organizations to secure IP communications, and their use is growing rapidly. While SSL provides data privacy and secure communications, it also creates challenges to inspection devices in the security stack when inspecting the encrypted traffic. In short, the encrypted communications cannot be seen as clear text and are passed through without inspection, becoming security blind spots. This creates serious risks for businesses: What if attackers are hiding malware inside the encrypted traffic?

However, performing decryption of SSL/TLS traffic on the security inspection devices, with native decryption support, can tremendously degrade the performance of those devices. This performance concern becomes even more challenging given the demands of stronger, 2048-bit certificates.

An integrated F5 and Cisco Advanced Malware Protection (AMP) solution solves these two SSL/TLS challenges. F5® SSL Orchestrator™ centralizes SSL inspection across complex security architectures, enabling flexible deployment options for decrypting and re-encrypting user traffic. It also provides intelligent traffic orchestration using dynamic service chaining and policy-based management. The decrypted traffic is then inspected by one or more Cisco next-generation firewalls (NGFWs), which can prevent previously hidden threats and block zero-day exploits. The Cisco Firepower Threat defense may be delivered using several combinations of Cisco Firepower and ASA platforms and software images. This solution eliminates the blind spots introduced by SSL and closes any opportunity for adversaries.

This guide provides an overview of the joint solution, describes different deployment modes with reference to service chain architectures, recommends practices, and offers guidance on how to handle enforcement of corporate Internet use policies.

The Integrated Solution

The F5 and Cisco integrated solution enables organizations to intelligently manage SSL while providing visibility into a key threat vector that attackers often use to exploit vulnerabilities, establish command and control channels, and steal data. Without SSL visibility, it is impossible to identify and prevent such threats at scale.

Key highlights of the joint solution include:

- **Flexible deployment modes** that easily integrate into even the most complex architectures, consolidate the security stack to reduce complexity, and deliver SSL visibility across the security infrastructure.
- **Centralized SSL decryption/re-encryption** with best-in-class SSL hardware acceleration, eliminating the processing burden of multiple decryption/re-encryption workloads on every security inspection hop in the stack, which reduces latency while improving the user experience.
- **Dynamic security service chaining**, which provides policy-based traffic management, thus determining whether traffic should be allowed to pass or be decrypted and sent through a security device or service.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

- **An industry-leading application delivery controller** that load balances traffic to multiple devices in the security services, enabling effortless scaling and growth.
- **Built-in health monitors** that detect security service failures and shifts, or bypasses, loads in real time to provide reliability and fault tolerance.
- **Full cipher support**, including support for the perfect forward secrecy (PFS) enabled ciphers, to ensure full traffic visibility.
- **Advanced sandboxing capabilities** to perform automated static and dynamic analysis, then uncover stealthy threats and help the security team to understand, prioritize, and block sophisticated attacks.
- **Point-in-time malware detection and blocking** using anti-virus (AV) detection engines, one-to-one signature matching, machine learning, and fuzzy fingerprinting.
- **Global threat intelligence sharing** by Cisco experts who analyze millions of malware samples and push that intelligence to AMP to correlate against this context-rich knowledge base, which enables it to proactively defend against known and emerging threats.

F5 SSL Orchestrator installs a decryption/clear text zone between the client and web server, creating an aggregation (and, conversely, disaggregation) visibility point for security services.

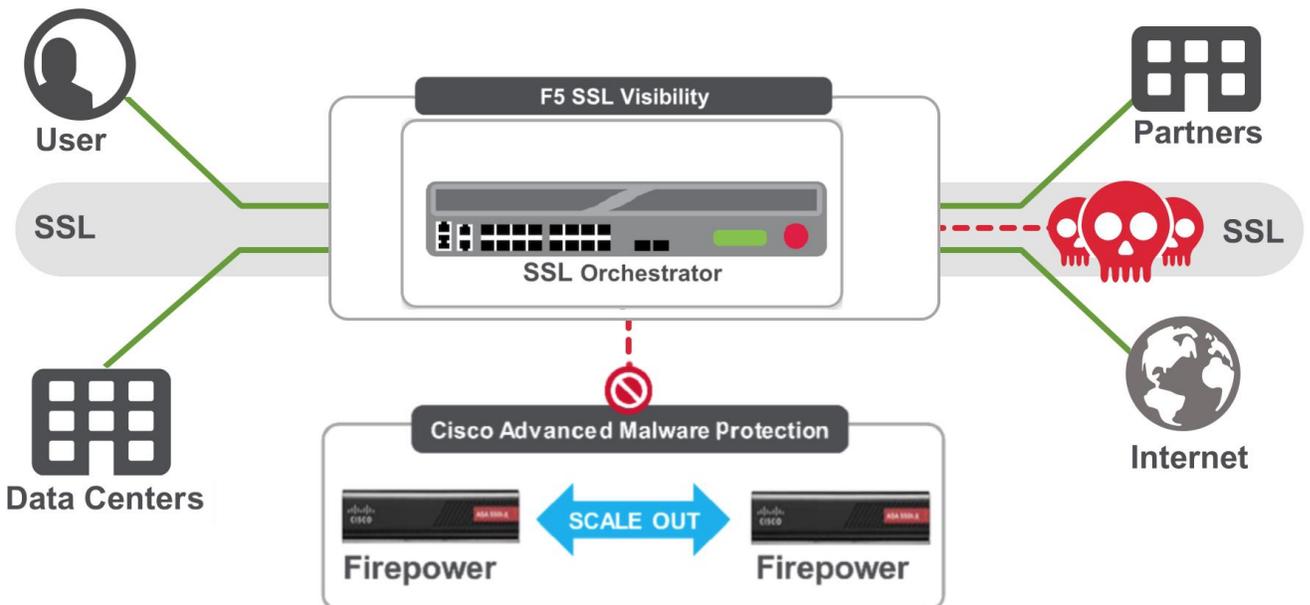


Figure 1: The integrated F5 and Cisco Firepower security solution

SSL visibility: How do we do it?

The F5 system establishes two independent SSL connections—one with the client and the other with the web server. When a client initiates an HTTPS connection to the web server, the F5 system intercepts and decrypts the client-encrypted traffic and steers it to a pool of Cisco Firepower devices (or devices running FirePOWER Services) for

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

inspection before re-encrypting the same traffic to the web server. The return HTTPS response from the web server to the client is likewise intercepted and decrypted for inspection before being sent on to the client.

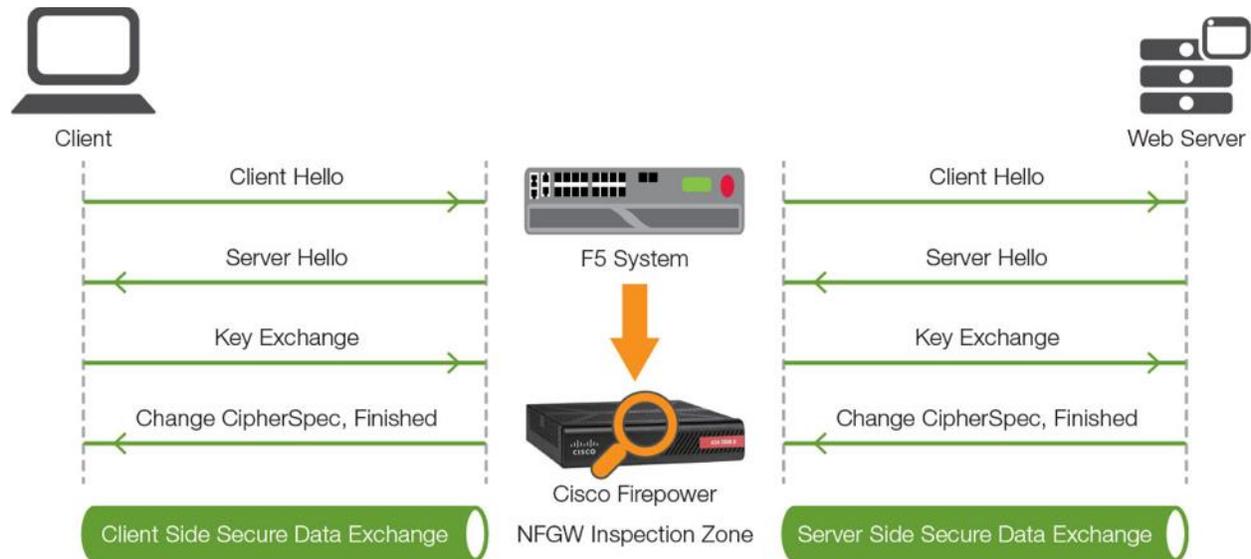


Figure 2: The F5 full proxy architecture

SSL orchestration using security service chains

A typical security stack often consists of more than advanced anti-malware protection systems. It begins with a firewall but almost never stops there, with components such as intrusion detection/prevention systems (IDS/IPS), web application firewalls, data loss prevention (DLP), and more. To solve specific security challenges, security administrators are accustomed to manually chaining these multiple point security products by creating a bare-bones security stack consisting of multiple services. In this model, all user sessions are provided the same level of security, as this “daisy chain” of services is hard-wired.

As shown in Figure 3, SSL Orchestrator can load balance, monitor, and dynamically chain security services, including next-gen firewalls, DLP, IDS/IPS, web application firewalls, and anti-virus/malware, by matching the user-defined policies to determine whether to bypass or decrypt and whether to send to one set of security services or another. This policy-based traffic steering capability allows for better utilization of the existing security services investment and helps to reduce administrative costs.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

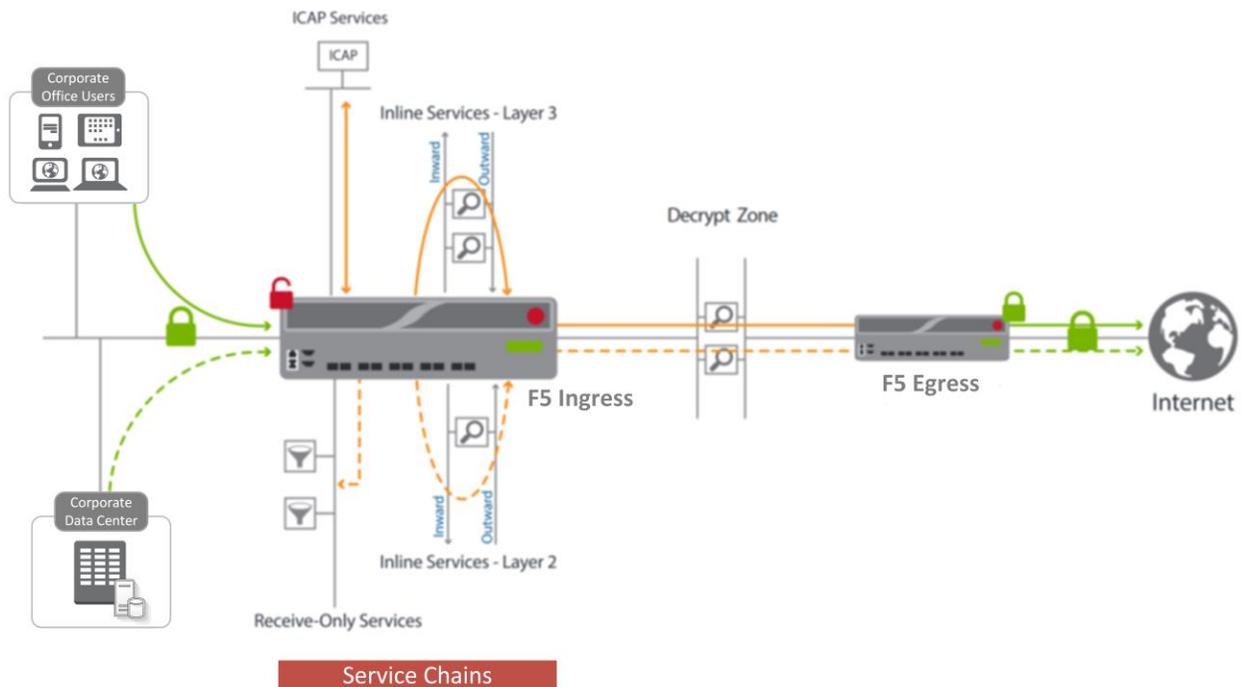


Figure 3: The security service chaining architecture

The F5 SSL visibility solution provides a way to apply different service chains based on context derived from a powerful classification engine. That context can come from:

- Source IP/subnet.
- Destination IP/subnet.
- IP intelligence category.
- IP geolocation.
- Host and domain name.
- URL filtering category.
- Destination port.
- Protocol.

Deployment Planning

Careful advance consideration of deployment options can ensure an efficient and effective implementation of the F5 integrated solution using the Cisco Firepower security system.

Sizing

The main advantage of deploying SSL Orchestrator in the corporate security architecture is that the wire traffic now can be classified as “interesting” traffic, which needs to be decrypted by SSL Orchestrator for inspection by Cisco

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

Firepower, and “uninteresting” traffic, which is allowed to pass through or be processed differently according to other corporate policy requirements. This selective steering of only the interesting traffic to the firewall system conserves its valuable resources (as it need not inspect the entire wire traffic), maximizing performance.

As a result, it is important to consider the entire wire traffic volume to calculate the appropriate F5 device size.

Depending on the [mode of deployment](#) you choose, you will need at least two interfaces on the F5 system for each firewall configured for inline mode and at least one interface for each firewall configured for TAP mode.

Refer to the [SSL Orchestrator Datasheet](#) and consider the following factors when sizing the F5 system for the integrated solution:

- Port density
- SSL bulk encryption throughput
- System resources
- The number of security services and devices in them

License components

F5 SSL Orchestrator hardware—the i2800, i5800, i10800—supports this integration. By default, SSL Orchestrator ships with an installed base module that provides both SSL interception and service chaining capabilities.

SSL Orchestrator can also be deployed as an application on an existing F5® BIG-IP® system. Please contact your local F5 representative to understand the licensing and deployment options. For simplicity’s sake, references to SSL Orchestrator and the BIG-IP system in this document (and some user interfaces) apply equally regardless of the F5 hardware used. The solution architecture and configuration are identical.

Optionally, customers can consider the following:

- **A URL Filtering (URLF) subscription** to use the URL category database for filtering.
- **An F5 IP Intelligence subscription** to detect and block known attackers and malicious traffic.
- **A network hardware security module (HSM)** to safeguard and manage digital keys for strong authentication.

Cisco Firepower can be deployed:

- Via Firepower Threat defense (a unified software image) on the ASA 5000x and Firepower 2100/4100/9300 platforms.
- Via FirePOWER services on a separate FirePOWER module on an ASA 5500x platform.

Horizontal scaling

SSL Orchestrator’s ability to steer and load-balance traffic to multiple security devices in a service or service pool enables the Cisco security platform to scale horizontally without the need for any functional add-on. This ensures that the service is not only fault-tolerant but also highly available, maximizing throughput.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

It is common to configure a single pool of Cisco Firepower NGFWs with SSL Orchestrator load balancing the unencrypted HTTP and decrypted HTTPS traffic to all the pool members. However, if you need to create multiple firewall pools, each taking a different traffic set based on user-defined criteria such as VLAN, tenant, or other criteria, you can do so by leveraging the TCP Service Chain Classifier Rules in SSL Orchestrator. These rules classify the wire traffic based on user-defined network information, IP geolocation, URL category, protocol, or IP intelligence, among other factors, and steer the classified traffic accordingly to a designated service chain the Cisco Firepower pool is part of.

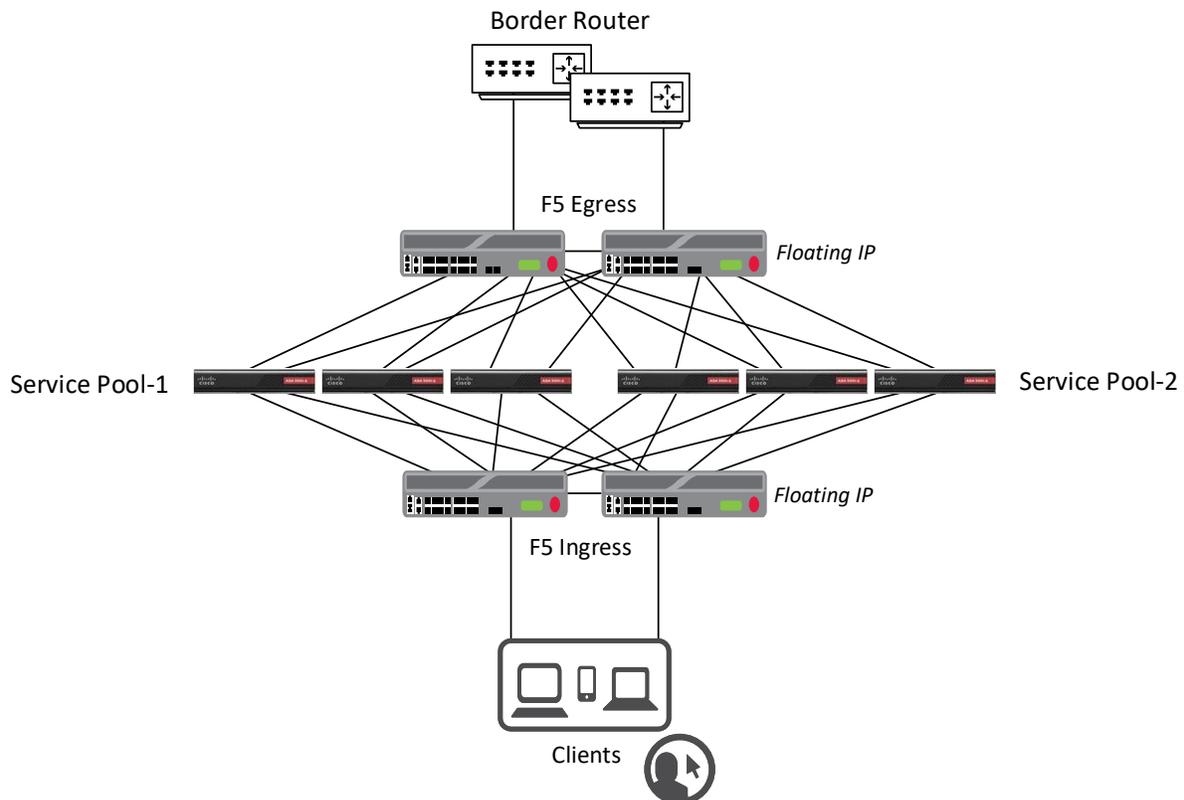


Figure 4: Cisco Firepower NGFWs in a service/service pool scaling configuration horizontal with SSL Orchestrator

Traffic exemptions for SSL inspection

As noted, SSL Orchestrator can be configured to distinguish between interesting and uninteresting traffic for the purposes of security processing. Examples of uninteresting traffic (including those types that cannot be decrypted) to be exempted from inspection may include:

- Guest VLANs.
- Applications that use pinned certificates.
- Trusted software update sources like Microsoft Windows updates.
- Trusted backup solutions like a crash plan.
- Any lateral encrypted traffic to internal services to be exempted.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

You can also exempt traffic based on domain names and URL categories. The service chain classifier rules of SSL Orchestrator enable administrators to enforce corporate Internet use policies, preserve privacy, and meet regulatory compliance.

Traffic exemptions based on URL category might include bypasses (and thus no decryption) for traffic from known sources of these types of traffic:

- Financial
- Health care
- Government services

Certificate requirements

An SSL certificate—preferably a subordinate certificate authority (CA)—and private key on SSL Orchestrator are needed to generate and issue certificates to the end host for client-requested HTTPS websites that are being intercepted.

To ensure that clients on the corporate network do not encounter certificate errors when accessing SSL-enabled websites from their browsers, the root certificate must be imported into the browser or operating system of the end hosts.

IP addressing

When a Cisco Firepower NGFW is deployed as an L3/routed hop, we recommend configuring its IP addresses for connected inward and outward VLANs from default fixed addressing subnets, provided by SSL Orchestrator, that are derived from a RFC2544 CIDR block of 192.19.0.0. This minimizes the likelihood of address collisions.

For example, you can configure a firewall to use the IP address 198.19.0.61/25 on the inward VLAN and 198.19.0.161/25 on the outward VLAN pointing to the F5 connected interfaces. You will also need to configure static routes to the internal networks on the firewall inward VLAN and a default route to the Internet on the outward VLAN. The table below explains the IP addresses that you need to configure when deploying multiple firewalls in the service pool.

Cisco Firepower NGFW	Inward Interface IP	Inward / Internal Gateway	Outward Interface IP	Outward/ Default Gateway
Cisco Firepower NGFW-1	198.19.0.61/25	198.19.0.10/25	198.19.0.161/25	198.19.0.245/25
Cisco Firepower NGFW-2	198.19.0.62/25		198.19.0.162/25	
Cisco Firepower NGFW-n $n \leq 8$	198.19.0.6n/25		198.19.0.16n/25 $n \leq 8$	

Deployment modes

Due to security concerns around key compromise, Internet sites have started to move away from RSA-based encryption. RSA, as a key exchange encryption protocol, uses the server's key pair to negotiate the symmetric keys used in the encrypted session, therefore potentially compromising the server's private key (such as in the Heartbleed vulnerability), as well as compromising any message, current or past, that uses or used that key pair. Therefore, these websites are transitioning to encryption technologies based on Diffie-Hellman (DH) key agreement protocols that do not expose data if a private key is compromised. Further, making DH keys ephemeral (temporary) defines that cryptography as perfect forward secrecy (PFS). PFS protects past sessions against future compromise of the secret keys, as they are not linked to the server's key pair.

An interesting side effect of this evolution is that passive SSL inspection technologies—systems that exist in the market today and can attach to a span port to passively (and often asynchronously) decrypt SSL/TLS communications—can no longer function. These technologies rely on the client and server performing an RSA key exchange, and they must possess a copy of the server's private key. If the client and server choose a PFS cipher, there is no opportunity for these passive SSL systems to decrypt the data. Many Internet sites and most browsers today prefer PFS ciphers over non-PFS (RSA) ciphers. In addition, the newest TLS version 1.3 update will completely remove non-PFS key exchanges, making passive SSL systems nonfunctional. In other words, to perform SSL visibility when employing ciphers based on PFS, an intercept system must be inline to the traffic flow.

Within that provision, various modes of deployment are available for integrating F5 systems with Cisco Firepower firewalls for advanced threat protection.

Single or double F5 systems

The F5 SSL visibility solution with inline Cisco Firepower NGFWs can be deployed with one or two F5 systems.

- **Option 1: A single F5 system with inline Cisco Firepower NGFW.** This solution entails a single F5 system deployed to perform both decryption and re-encryption of SSL traffic, while Cisco Firepower NGFWs are configured for inline mode and deployed as an L3 service pool on the F5 system.
- **Option 2: Two F5 systems with inline Cisco Firepower NGFW.** Although a single F5 system delivers all the capabilities and functionality needed to deploy the SSL visibility solution, in some cases, customers may want to implement a second F5 system on the egress:
 - When there is a need for increased SSL throughput of the solution, or
 - When the organization's security policy dictates deploying two SSL intercept appliances for visibility.

When using two F5 systems for the SSL visibility solution, the ingress system on the client side will decrypt the client-encrypted web traffic, while the egress system on the server/Internet side will re-encrypt the same traffic before sending it to the web server, maximizing SSL throughput.

Service pool or sandwich

Deploying the SSL visibility solution using two F5 systems entails two options for configuring the Cisco Firepower NGFWs:

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

- **As a service pool managed on the ingress F5 system.** The advantage of deploying firewalls in a service pool is that the ingress F5 system can then steer traffic based on user-defined service chain policies.
- **With the pool sandwiched between the ingress and egress F5 systems.** While this type of deployment allows the F5 system to steer unencrypted traffic that has traveled through the service chains to the sandwiched firewalls in the decrypt zone, it has the downside of limiting policy enforcement, as the firewall pool is no longer a part of any services chain. Further, it also complicates the failover design. The firewalls now need to be configured to fail open, as the decrypt zone has no built-in way to go around device failures here.

Both of these mode options are valid for outbound flows (for example, corporate users browsing the web over HTTPS). They are also applicable at any data exchange points in the data center where the encrypted traffic flows outbound from one security zone to another.

Architecture best practices

A number of best practices can help ensure a streamlined architecture that optimizes performance and reliability as well as security. F5 recommendations include:

- Deploy the F5 systems in a [sync/failover device group](#) (S/FDG), which includes the active-standby pair, with a floating IP address for high availability (HA).
- Every Cisco Firepower NGFW in the service pool must be dual-homed on the inward and outward VLANs with each F5 system in the device sync/failover device group.
- Further interface redundancy can be achieved using the Link Aggregation Control Protocol (LACP). LACP manages the connected physical interfaces as a single virtual interface (aggregate group) and detects any interface failures within the group.
- Unlike with some competing solutions, the F5 systems do not need physical connections to the Cisco Firepower NGFWs. All the F5 system requires is L3 reachability to steer traffic through the firewalls. In slow networks, however, we recommend deploying the firewalls not more than one hop away. As a generic guideline, when inspection devices are not directly connected to the F5 system, we highly recommend use of network and VLAN controls to restrict access to the unencrypted data only to the inspection devices. For the same reason, RFC2544 addresses mandated by SSL Orchestrator provide that extra level of security control at the network layer.

Initial Setup

Initial setup includes configuration of Cisco Firepower on ASA and setup of SSL Orchestrator. Once these steps are complete, you can proceed to configuration for the specific deployment scenario you choose.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

- ii. Under **User Administration**, enter and confirm the **Root Account** and **Admin Account** passwords, and click **Next**. The Root Account provides access to the command line, while the Admin Account accesses the user interface.

The screenshot displays the F5 Hercules configuration interface. At the top, it shows the status 'ONLINE (ACTIVE) Standalone' and the 'HERCULON' logo. The navigation menu includes 'Main', 'Help', and 'About'. The left sidebar contains icons for 'SSL Orchestrator', 'Statistics', 'Local Traffic', 'Device Management', 'Network', and 'System'. The main content area is divided into two sections: 'General Properties' and 'User Administration'.
General Properties:
- Management Port Configuration: Automatic (DHCP) Manual
- Management Port: IP Address/prefix: 192.168.16.31, Network Mask: 255.255.255.0, Management Route: 192.168.16.10
- Host Name: SSLO.f5sec.com
- Host IP Address: Use Management Port IP Address
- Time Zone: America/Los Angeles
User Administration:
- Root Account: Disable login, Password: [masked], Confirm: [masked]
- Admin Account: Password: [masked], Confirm: [masked]
- SSH Access: Enabled
- SSH IP Allow: * All Addresses
At the bottom, there are 'Back' and 'Next...' buttons.

Figure 6: Platform configuration

7. The system notifies you to log out and then log back in with your username (*admin*) and new password. Click **OK**. The system reboots.
8. Once the **Network Time Protocol (NTP)** configuration screen opens, enter the IP **Address** of the NTP server to synchronize the system clock with, and click **Add**. Click **Next**.
9. (Optional, unless you plan to later use the DNSSEC option in the SSL Orchestrator configuration—in which case this step is required.) The **Domain Name Server (DNS)** screen opens. Complete the following steps:
 - i. To resolve host names on the system, set up the DNS and associated servers: For the **DNS Lookup Server List**, type the **IP Address** of the DNS server and click **Add**.
 - ii. If you use BIND servers, add them in the **BIND Forwarder Server** list.
 - iii. Add local domain lookups (to resolve local host names) in the **DNS Search Domain** list.
 - iv. Click **Next**. The **Internal VLAN** screen opens.
10. On the **Internal VLAN** screen, specify the **Self IP** settings for the internal network:
 - i. Enter a self IP **Address**.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

- ii. Enter a network mask (**Netmask**) for the self IP address.
- iii. Retain the default values for the **Port Lockdown** and **VLAN Tag ID** settings.
- iv. Under **Interfaces**, select an interface number from the **VLAN Interfaces** list, and then select Tagged or Untagged from the **Tagging** list. (Select **Tagged** when you want traffic for that interface to be tagged with a VLAN ID.) Click **Add**.
- v. Click **Next**. This completes the configuration of the internal VLAN.

The screenshot shows two configuration screens. The top screen, 'Internal Network Configuration', has a 'Select VLAN' dropdown set to 'internal'. Under 'Self IP', the 'Address' is '10.10.10.10', 'Netmask' is '255.255.255.0', and 'Port Lockdown' is 'Allow Default'. The bottom screen, 'Internal VLAN Configuration', has 'VLAN Name' set to 'internal' and 'VLAN Tag ID' as an empty field. Under 'Interfaces', 'VLAN Interfaces' is '2.0' and 'Tagging' is 'Untagged'. An 'Add' button is visible, and a table below it contains one entry: '1.0 (untagged)'. 'Edit' and 'Delete' buttons are at the bottom of the table. At the very bottom of the form are 'Back' and 'Next...' buttons.

Figure 7: Internal VLAN configuration

11. The **External VLAN** screen opens. Specify the **Self IP** settings for the external network:
 - i. Enter a self IP **Address**.
 - ii. Enter a network mask (**Netmask**) for the self IP address.
 - iii. Retain the default value for the **Port Lockdown** setting.
 - iv. Enter the IP address you want to use as the **Default Gateway** to the external VLAN.
 - v. Retain the default value (auto) for the **VLAN Tag ID** setting. Click **Next**. This completes the configuration of the external self IP addresses and VLAN.
12. On the **Forward Proxy Certificate** screen, complete the following configuration to import the CA certificate:
 - i. For the **Certificate Name**, select **Create New** and enter a name.
 - ii. For the **Certificate Source**, either select **Upload File** and choose a file, or select **Paste Text** and use copy and paste to enter your certificate source.
 - iii. For the **Key Source**, either select **Upload File** and choose a file, or select **Paste Text** and use copy and paste to enter your key source.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

- iv. If your certificate/key source is protected by a passphrase, select **Password** as the **Security Type**, and enter the passphrase. Otherwise leave the default setting. Click **Next**.
13. On the **Logging** screen, select either local or Splunk as the **Publisher Type**.
 - If you select local, specify the **Destination**—either local-db or localsyslog. This determines the destination of your logs, either a local database or a localsyslog server.
 - If you select Splunk, for **Protocol**, select either TCP or UDP. Enter the **IP Address** and **Port** of the Splunk server.
14. Click **Finish**. The SSL Orchestrator configuration page appears with a complete menu displayed on the left side of the page.

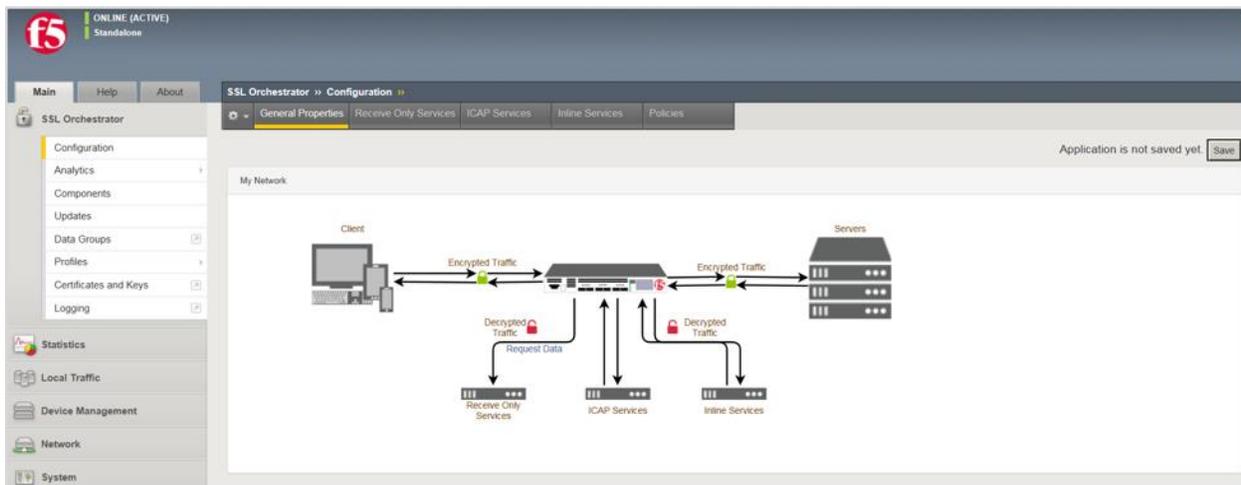


Figure 8: The SSL Orchestrator configuration screen once the initial setup is complete

You are now ready to proceed to the second part of configuration, where you finalize your system for SSL Orchestrator.

Update the SSL Orchestrator version

Periodic updates are available for the SSL Orchestrator configuration utility. To download the latest, follow these steps:

1. Visit downloads.f5.com. You will need your registered F5 credentials to log in.
2. Click **Find a Download**.
3. Scroll to the **Security** product family and select **SSL Orchestrator**.

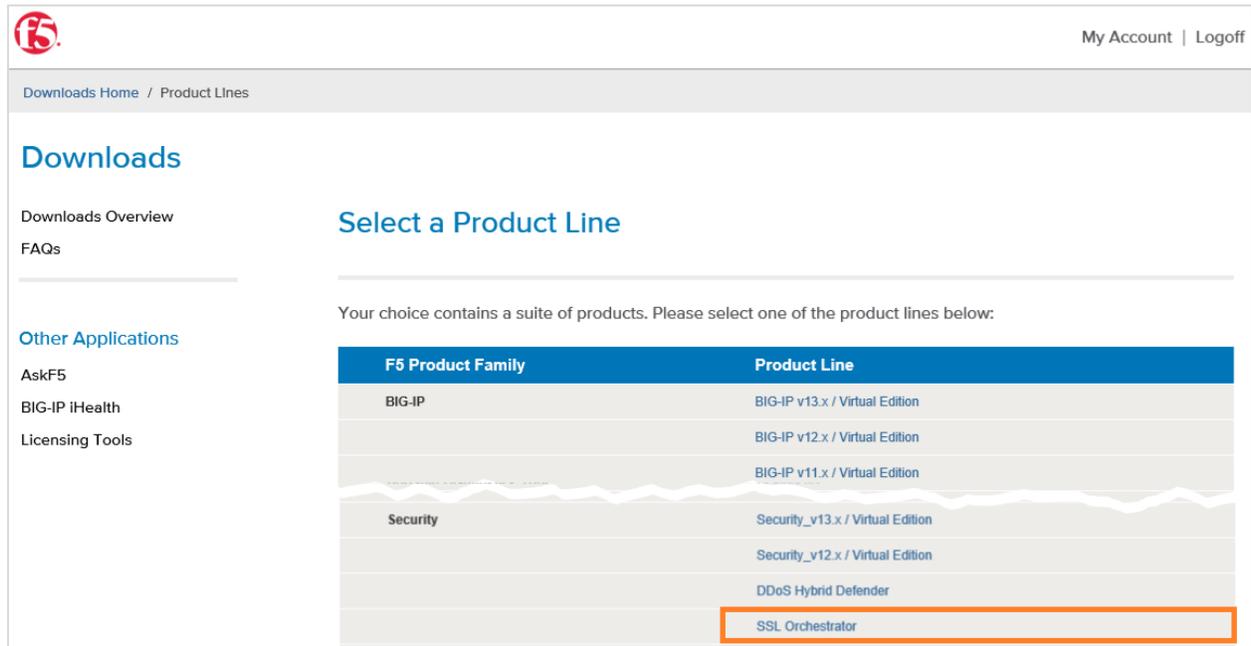


Figure 9: F5 product download web page

4. Click the SSL Orchestrator container.
5. Select and download the latest version of the SSL Orchestrator .rpm file.
6. Read through the appropriate [Release Notes](#) before attempting to use the downloaded file.
7. Once you've read the release notes, log in to the main tab of the F5 system management interface and navigate to **SSL Orchestrator > Updates**.
8. Under **File Name**, click **Browse** and navigate to the .rpm file you saved on your system. Click **Open** to select it.



Figure 10: Updating SSL Orchestrator

9. Click **Install**. The latest version of the SSL Orchestrator configuration utility will be installed. Your system may reboot to make the change effective.

Back up your F5 system configuration

Before beginning the detailed SSL Orchestrator configuration, we strongly recommend you back up the F5 system configuration using the following steps. This enables you to restore the previous configuration in case of any issues.

1. From the main tab of the F5 management interface, click **System > Archives**.
2. To initiate the process of creating a new UCS archive (backup), click **Create**.
3. Enter a *unique File Name* for the backup file.
4. Optional:
 - If you want to encrypt the UCS archive file, from the **Encryption** menu, select **Enabled** and enter a passphrase. You must supply the passphrase to restore the encrypted UCS archive file.
 - If you want to exclude SSL private keys from the UCS archive, from the **Private Keys** menu, select **Exclude**.



General Properties	
File Name	SSLO-state0
Encryption	Disabled
Private Keys	Include
Version	BIG-IP 13.0.0 Build 0.0.1645

Cancel Finished

Figure 11: New system archive creation

5. Click **Finished** to create the UCS archive file.
6. When the backup process is done, examine the status page for any reported errors before proceeding to the next step.
7. Click **OK** to return to the Archive List page.
8. Copy the .ucs file to another system.

To restore the configuration from a UCS archive, navigate to **System > Archives**. Select the name of the UCS file you want to restore and click **Restore**. For details and other considerations for backing up and restoring the F5 system configuration, see Solution K13132 on AskF5: [Backing up and restoring configuration files](#).

Configuration for a Single F5 System with FirePOWER Services on Cisco ASA in L2 Mode (Burrito Design)

This deployment mode entails a single F5 system performing SSL visibility. This single system handles both decryption and re-encryption of HTTPS traffic, with an inspection zone installed between the ingress and the egress.

Figure 12 shows a standalone F5 system configured to intercept, decrypt, and steer the decrypted traffic to a service pool of two Cisco FirePOWER Services modules on ASAs configured in L2 mode where the traffic will be inspected for hidden threats. You can also deploy the F5 system as a device sync/failover device group (including an HA pair) with a floating IP address for high availability.

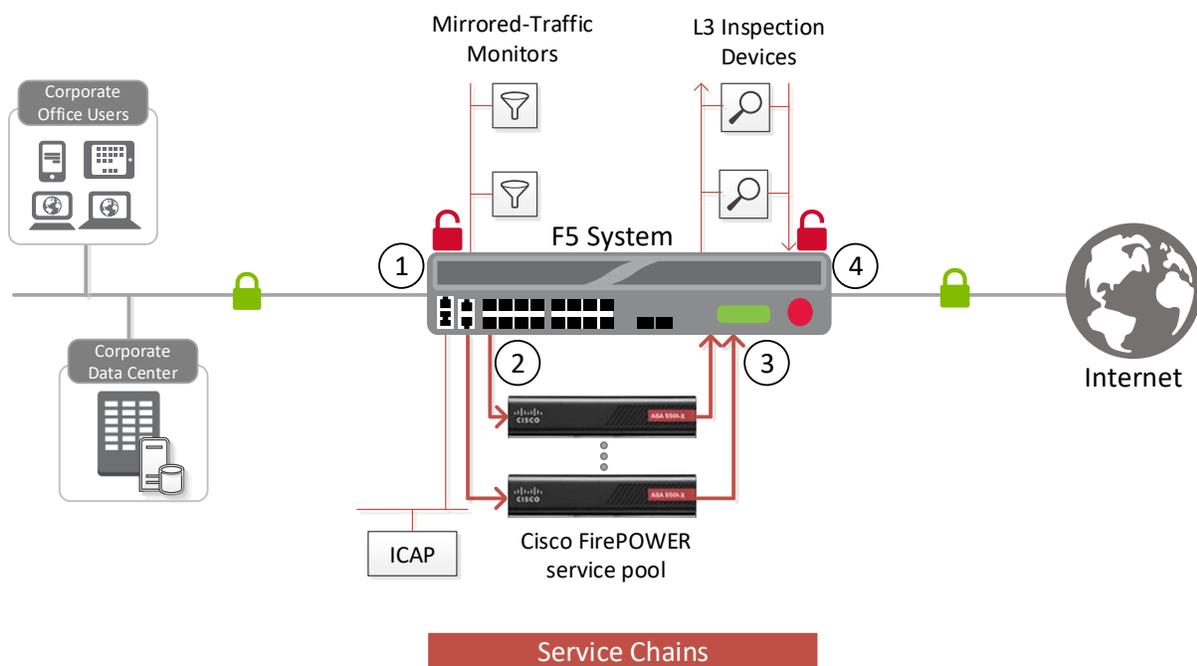


Figure 12: The SSL visibility solution with a pool of Cisco ASAs in a service chain using one F5 system

How traffic flows in this deployment option:

1. Client traffic arriving at the ingress F5 device is classified, and interesting HTTPS traffic is decrypted as part of the SSL forward proxy process.
2. The ingress virtual servers on the F5 system steer the decrypted traffic through a service pool of Cisco ASAs as part of a service chain via a service-inward VLAN.
3. The HTTP traffic is inspected by the FirePOWER Services for any hidden threats before sending that traffic back to the F5 system on the service-outward VLAN.

- The F5 system orchestrates the decrypted traffic through other services in the chain before it aggregates and re-encrypts the traffic, which is routed outbound to the web server.

Configure SSL Orchestrator

In the example configuration below, SSL Orchestrator steers the outbound web traffic through the Cisco ASAs, which are part of a service chain of security devices. Please refer to the [F5 SSL Orchestrator Setup Guide](#) for additional help during configuration.

General properties

This first step must be completed before you can set up services, service chains, and classifier rules.

- On the main screen of the management console, click **SSL Orchestrator > Configuration > General Properties**.
- Answer the configuration questions (see Figure 13) for SSL Orchestrator. (Also see the table below for examples and tips.)

Question	User Input
Application Service Name	Enter a name without spaces or dashes for the SSL Orchestrator application.
Do you want to set up separate ingress and egress devices with a cleartext zone between them?	<p>You can configure a single F5 device to receive both ingress and egress traffic on different networks, or you can configure separate F5 devices for ingress and egress traffic. If you choose the latter option, you are asked further questions to enter peer application names, control channel virtual server IPs, and pre-shared keys to establish and protect the communication between the devices.</p> <p>Otherwise, select No, use one BIG-IP device for ingress and egress. (Select this option even if you are using new SSL Orchestrator hardware.) This sample configuration follows that option.</p>
Which IP address families do you want to support?	Select Support IPv4 only . (Currently SSL Orchestrator only supports IPv4 families.)
Which proxy schemes do you want to implement?	<p>SSL Orchestrator can operate in transparent and/or explicit proxy mode. If you choose explicit proxy, a separate explicit proxy configuration section displays for you to choose the VLANs that explicit proxy needs to listen to and so you can enter the IP address and port number of the explicit proxy.</p> <p>Select Implement Transparent proxy only.</p>
Do you want to pass UDP traffic through the transparent proxy unexamined?	<p>This option only applies if you selected Implement transparent proxy only above. By default, transparent proxy mode manages TCP traffic but allows UDP traffic to pass through unexamined. Choose No to prevent the passage of unexamined UDP traffic.</p> <p>Otherwise, select the default, Yes, pass all UDP traffic unexamined.</p>
Do you want to pass non-TCP, non-UDP traffic through the transparent proxy?	<p>This option also only applies if you select Implement transparent proxy only. By default, transparent proxy mode passes through non-TCP, non-UDP traffic (such as IPSec, SCTP, and OSPF). Choose No to block.</p> <p>Otherwise, select the default, Yes, pass non-TCP, non-UDP traffic.</p>

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

Which is the SSL Forward Proxy CA certificate?	Select the Certificate Authority (CA) certificate that your clients will trust to authenticate intercepted TLS connections. If you did not use the Setup Wizard, you must import a CA certificate before you can use this functionality.
Which is the SSL Forward Proxy CA private key?	Select the corresponding private key, which you imported with the CA certificate while configuring the Setup Wizard. If you did not use the Setup Wizard, you must import a CA certificate before you can use this functionality.
What is the private-key passphrase (if any)?	Enter the private-key passphrase, if any. If the key does not have a passphrase, leave this field empty.
Which CA bundle is used to validate remote server certificates?	<p>The CA bundle is the collection of root and intermediate certificates for the CA you trust to authenticate servers where your clients might connect. The CA bundle is also known as the local trust store.</p> <p>Select the CA bundle that validates the remote server certificates.</p>
Should connections to servers with expired certificates be allowed?	<p>Remote servers can present expired certificates. Allowing connections to servers with expired certificates can cause a security risk. Legitimate servers do sometimes offer certificates which are overdue for renewal or which were signed by legitimate CAs but that are simply unknown to the F5 system. In the latter case, if you allow connections, consider adding any needed CA certificates to the F5 system CA bundle (trust store).</p> <p>Select No, forbid connections to servers with expired certificates to prevent connections to servers that have expired certificates.</p>
Should connections to servers with untrusted certificates be allowed?	<p>Remote servers can present untrusted certificates. Allowing connections to servers with untrusted certificates can cause a security risk.</p> <p>Select Yes, allow connections to servers with untrusted certificates if appropriate for your situation and security policies.</p>
Should strict updates be enforced for this application?	<p>If you select this option, you cannot manually modify any settings produced by the application. Once you disable this option, you can manually change your configuration.</p> <p>F5 recommends enabling this setting (select Yes) to avoid misconfigurations that can cause an unusable application.</p>

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

General Properties	
Application Service Name [?]	SSLVisibility
Do you want to setup separate ingress and egress devices with a cleartext zone between them? [?]	No, use one BIG-IP device for ingress and egress
Which IP address families do you want to support? [?]	Support IPv4 only
Which proxy schemes do you want to implement? [?]	Implement transparent proxy only
Do you want to pass UDP traffic through the transparent proxy unexamined? [?]	Yes, pass all UDP traffic unexamined
Do you want to pass non-TCP, non-UDP traffic through the transparent proxy? [?]	Yes, pass non-TCP, non-UDP traffic
Which is the SSL Forward Proxy CA certificate? [?]	/Common/Sub-CA.crt
Which is the SSL Forward Proxy CA private key? [?]	/Common/Sub-CA.key
What is the private-key passphrase (if any)? [?]	
Which CA bundle is used to validate remote server certificates? [?]	/Common/ca-bundle.crt
Should connections to servers with expired certificates be allowed? [?]	No, forbid connections to servers with expired certificates
Should connections to servers with untrusted certificates be allowed? [?]	No, forbid connections to servers with untrusted certificates
Should strict updates be enforced for this application? [?]	<input checked="" type="checkbox"/> Enabling strict updates enforces protection of your configuration by restricting the ability to modify objects outside of this application.

Figure 13: Sample general properties configuration

- Continue configuration by scrolling down to **Ingress Device Configuration** (see below.)

Ingress device configuration

The ingress device is one or more ingress VLANs where clients send traffic. The F5 device decrypts the encrypted traffic on ingress and then, based on protocol, source, and destination, classifies the traffic and passes each connection for inspection.

- Answer each configuration question. See tips and guidance below.

Question	User Input
Which VLAN(s) will bring client traffic to the transparent proxy?	Select one or more VLANs where transparent-proxy ingress traffic will arrive.
How should a server TLS handshake failure be handled?	Most TLS handshake failures occur during protocol and cipher agreement. You can specify whether to drop or bypass the connection. Typically, select If server TLS handshake fails the connector fails .
DNS query resolution	Specify whether to permit the system to send DNS queries directly to the Internet, or specify one or more local forwarding nameservers to process all DNS queries from SSL Orchestrator. If you choose the former, you can specify to configure local/private DNS zones. In this example, select Send DNS queries to forwarding nameservers on the local network .

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

Which local forwarding nameserver(s) will resolve DNS queries from this solution?	Type the IP address of the local nameserver(s) which will resolve the DNS queries.
Do you want to use DNSSEC to validate DNS information?	DNSSEC is a suite of extensions that add security to the DNS protocol by enabling DNS responses to be validated. Select Yes, use DNSSEC to validate DNS information.

Ingress Device Configuration

Which VLAN(s) will bring client traffic to the transparent proxy? ?	<div style="display: flex; justify-content: space-between;"> <div style="border: 1px solid #ccc; padding: 5px; width: 45%;"> <p>Selected</p> <p>Filter</p> <p>/Common/Internal</p> </div> <div style="border: 1px solid #ccc; padding: 5px; width: 45%;"> <p>Available</p> <p>/Common/External</p> </div> </div>
How should a server TLS handshake failure be handled? ?	If server TLS handshake fails then connector fails v
DNS query resolution ?	Send DNS queries to forwarding nameservers on the local network v
Which local forwarding nameserver(s) will resolve DNS queries from this solution? ?	<p>Nameserver IP address: <input style="width: 100px;" type="text"/> Add</p> <p>192.168.16.10</p> <p><input style="width: 100px;" type="text"/></p> <p>Delete</p>
Do you want to use DNSSEC to validate DNS information? ?	Yes, use DNSSEC to validate DNS information v

Figure 14: Sample ingress device configuration

- Continue configuration by scrolling down to **Egress Device Configuration** (see below.)

Egress device configuration

The egress device is one or more egress VLANs where the clients receive traffic. The F5 system decrypts the encrypted response on egress and then, based on protocol, source, and destination, classifies the traffic and passes each connection for inspection before sending it to the requested internal client.

- Answer each configuration question. Note that in this example, the same F5 device is configured to receive both the ingress and egress traffic.

Question	User Input
Do you want to SNAT client IP addresses?	It is common to translate the client source IP address with the address belonging to the egress for outbound traffic. Choose No to preserve the client source IP address. Otherwise, select Yes, SNAT (replace) client addresses
Do you want to use a SNAT Pool?	F5 recommends use of a SNAT pool to scale translations instead of overloading the egress interface IP address (AutoMap). Select Yes, define SNAT Pool addresses for good performance.
IPv4 SNAT addresses	Enter the IPv4 addresses for the SNAT pool.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

Should traffic go to the Internet via specific gateways?	Specify whether to route outbound using the default route on the F5 system or enter the IP address to be used as the default gateway. In the example above, we selected No, send outbound / Internet traffic via the default route .
---	--

Egress Device Configuration							
Do you want to SNAT client IP addresses? <small>?</small>	Yes, SNAT (replace) client addresses <input type="button" value="v"/>						
Do you want to use a SNAT Pool? <small>?</small>	Yes, define SNAT Pool addresses for good performance <input type="button" value="v"/>						
IPv4 SNAT addresses <small>?</small>	<table border="1"><thead><tr><th colspan="2">Address</th></tr></thead><tbody><tr><td>192.168.16.101</td><td><input type="button" value="+"/> <input type="button" value="-"/></td></tr><tr><td>192.168.16.102</td><td><input type="button" value="+"/> <input type="button" value="-"/></td></tr></tbody></table>	Address		192.168.16.101	<input type="button" value="+"/> <input type="button" value="-"/>	192.168.16.102	<input type="button" value="+"/> <input type="button" value="-"/>
Address							
192.168.16.101	<input type="button" value="+"/> <input type="button" value="-"/>						
192.168.16.102	<input type="button" value="+"/> <input type="button" value="-"/>						
Should traffic go to the Internet via specific gateways? <small>?</small>	No, send outbound / Internet traffic via the default route <input type="button" value="v"/>						

Figure 15: Sample egress device configuration

2. Continue configuration by scrolling down to **Logging Configuration** (see below.)

Logging configuration

1. Answer the configuration questions using the guidance below.

Question	User Input
What SSL Intercept logging level do you want to enable?	F5 recommends leaving the logging level at the default, Errors. Log on functional errors , unless you need to troubleshoot.
Which Log Publisher will process the log messages?	Specify whether to process the logs with an existing log publisher or that logs should be sent to syslog-ng.
What kind of statistics do you want to record?	Specify the kind of statistics you want the system to record. SSL Orchestrator can collect usage data for connections, service chains, services, and more. For optimal performance, keep the settings at the default, Usage counters only .

Logging Configuration	
What SSL Intercept logging level do you want to enable? <small>?</small>	Errors. Log on functional errors <input type="button" value="v"/>
Which Log Publisher will process the log messages? <small>?</small>	None (Send log messages to syslog-ng) <input type="button" value="v"/>
What kind of statistics do you want to record? <small>?</small>	Usage counters only (No remote-domain+cipher records) <input type="button" value="v"/>

Figure 16: Default logging settings

2. When you're done, click **Save** at the top of the page.

Create layer 2 inline service

Note: Before creating inline services, you must complete configuration of all the sections in the [General Properties](#) tab.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

Inline services pass traffic through one or more firewalls at layer 2 or layer 3. The firewalls are configured to communicate with the F5 system via two VLANs. In this section, we will configure an L2 inline service for a pair of Cisco ASAs.

1. On the main tab, click **SSL Orchestrator > Configuration > Inline Services**. The Inline Services screen displays.
2. Click **Add**.
3. Enter information for the configurable fields, following the guidance below.

Configuration Field	User Input
Name	Type a Name for the inline service.
Service Type	Select Layer 2 from the Service Type list.
Interfaces	Select the From BIG-IP and To BIG-IP system interfaces connected to the firewall and their respective VLANs Tags . Click Add . If you have multiple firewalls in the service pool, choose the F5 system interfaces connected to each firewall and their VLANs tags and click Add before moving to the next one. If you choose to use the Ratio field, the F5 system distributes connections among pool members in a static rotation according to ratio weights that you define. For example, if you have two devices, and one handles twice as much traffic as the other, you can set the ratio to 1 on the smaller device and 2 on the larger one.
Translate port for HTTP traffic	Select Use No if the connections should use their original destination ports . Choose Yes to translate the port for HTTP traffic to either 80, 8080, or 8443.
Connection Handling On Outage	Select Skip Service to allow connections to skip the service you are configuring if all the devices in the service pool are unavailable. Or select Reject Connection to reject every connection reaching the service when the service is down.

The screenshot shows the configuration page for an inline service. At the top, there are tabs for Name, Service Type, Interfaces, Available Devices, Translate Port for HTTP Traffic, and Connection Handling On Outage. The Name field contains 'Cisco'. The Service Type dropdown is set to 'Layer 2'. Under the Interfaces section, there are two columns: 'From BIG-IP' and 'To BIG-IP'. The 'From BIG-IP' column has a 'Ratio' field with '1', an 'Interface' dropdown with '1.3', and a 'Tag' field. The 'To BIG-IP' column has an 'Interface' dropdown with '1.4' and a 'Tag' field. Below these is an 'Add' button. To the right, there is an 'IP Address' field with '198.19.0.61' and an 'Add' button. At the bottom left, there are 'Finished' and 'Cancel' buttons. On the right side, there are dropdowns for 'Translate Port for HTTP Traffic' (set to 'No') and 'Connection Handling On Outage' (set to 'Skip Service').

Figure 17: Inline layer 2 service configuration

4. Leave other inline services configurations at their default settings.
5. When done, click **Finished**, then click **Save** at the top of the page.

Configuration for an F5 System with FirePOWER Services on Cisco ASA in L3 Mode

This deployment is similar to the solution explained above in the section called, "[Configuration for a Single F5 System with FirePOWER Services on Cisco ASA in L2 Mode.](#)" The only difference is that the inline service type is configured as an L3 service.

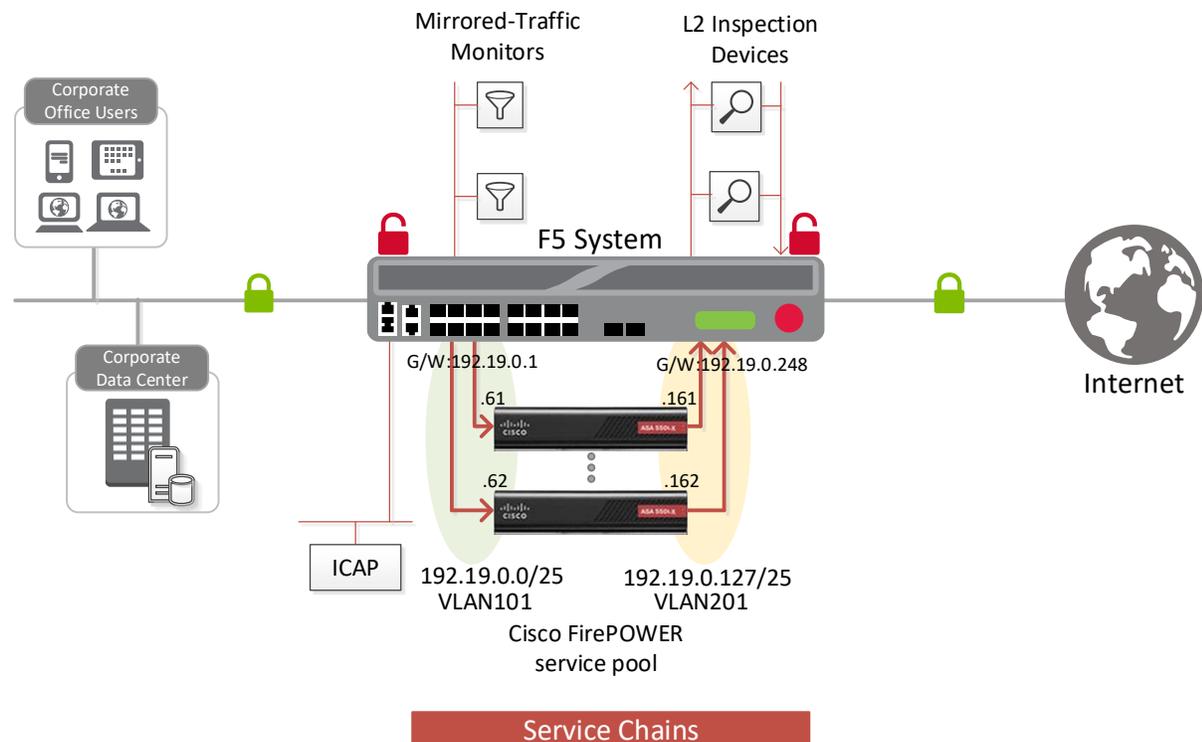


Figure 18: The SSL visibility solution with Cisco ASAs configured in L3 mode

Create the layer 3 inline service

The layer 3 inline service, as a prerequisite, requires you to configure the IP addresses on the firewall(s) from a specific fixed addressing scheme as explained in the section on [IP addressing](#). This configuration enables the F5 system to send to, and receive from, the firewall using a pre-defined set of addresses.

1. On the main tab, click **SSL Orchestrator > Configuration > Inline Services**.
2. Enter information for the configurable fields, following the guidance below.

Configuration Field	User Input
What is the IPv4 (CIDR /19) subnet-block base address?	For IPv4, F5 recommends the default block 198.19.0.0/19. Click Add . Even though you can change the base address of each address block (IPv4) from which subnets and addresses are assigned, changing an address block has several implications, must be done with caution, and is not recommended or supported by F5.
Name	Enter a Name for the inline service.
Service Type	Select Layer 3 from the Service Type list.
Interfaces	Select From BIG-IP and To BIG-IP system interfaces connected to the firewall(s) and their respective VLANs Tags .
Available Devices	Select the IP Address(es) of the firewall(s). In the sample configuration in Figure 19, the two configured IP addresses on each of the firewalls in the service pool are selected. These firewalls are preconfigured from the CIDR block of 192.19.0.0.
Translate port for HTTP traffic	Select No if the connections should use their original destination ports. Choose Yes to translate the port for HTTP traffic to either 80, 8080, or 8443.
Connection Handling On Outage	Select Skip Service to allow connections to skip the service you are configuring if all the devices in the service pool are unavailable. Or select Reject Connection for the system to reject every connection reaching the service when the service is down.

The screenshot displays the configuration page for an inline layer 3 service. At the top, there are tabs for Name, Service Type, Interfaces, Available Devices, Translate Port for HTTP Traffic, and Connection Handling On Outage. The main configuration area includes a Name field with 'Cisco', a Service Type dropdown set to 'Layer 3', and buttons for 'Finished' and 'Cancel'. The Interfaces section has 'From BIG-IP' and 'To BIG-IP' sections, each with an 'Interface' dropdown (set to 1.3 and 1.4) and a 'Tag' field. The Available Devices section shows an 'IP Address' list with '198.19.0.61' and '198.19.0.62', and an 'Add' button with '198.19.0.62' selected. The Translate Port for HTTP Traffic dropdown is set to 'No', and the Connection Handling On Outage dropdown is set to 'Skip Service'.

Figure 19: Inline layer 3 service configuration

- When done, click **Finished**, then click **Save** at the top of the page.

Configuration for the F5 System with Cisco ASAs in TAP Mode

In this solution option, the F5 system is configured to provide a packet-by-packet copy of both the unencrypted HTTP and decrypted HTTPS traffic to Cisco ASAs wherein the Cisco FirePOWER Services are configured for TAP mode.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

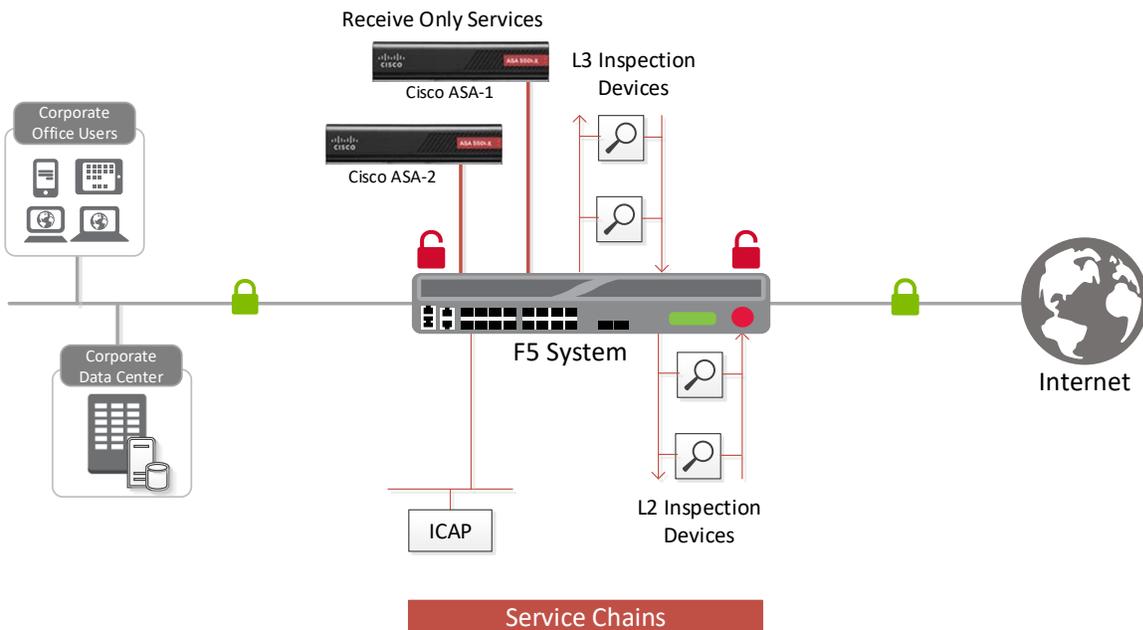


Figure 20: The SSL visibility solution for Cisco ASAs in TAP mode

Create a receive-only service

When firewalls are configured for receive-only service, they only receive traffic for inspection and do not send the traffic back to the F5 system. You can configure up to 10 receive-only services using the SSL Orchestrator configuration utility.

1. On the main tab, click **SSL Orchestrator > Configuration > Receive Only Services**.
2. Enter information for the configurable fields, following the guidance below.

Configuration Field	User Input
Name	Enter a Name for the receive-only service.
MAC Address	Enter the receiving interface's MAC Address . The MAC address can be obtained from the web UI.
IP Address	Enter the nominal IP Address for this device. Each receive-only device requires a nominal IP host address to identify the device in the F5 system. This nominal IP address must be homed on the same subnet as one (any one) of the F5 system's self-IP addresses. It does not have to be on the same VLAN as the receive-only device. No IP packets will ever be sent to the nominal IP address (but it must be unique on the network while it is assigned in this solution).
VLAN	From the VLAN list, select the VLAN where the receive-only device resides.
Interface	Select the associated F5 system Interface .

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

Receive Only Services ?				
Add Delete				
Name	MAC Address	IP Address	VLAN	Interface
<input type="checkbox"/> TAP	b4:0c:25:18:3e:10	192.168.16.100	/Common/External	1.1
Finished Cancel				

Figure 21: Sample receive-only service configuration

- When done, click **Finished**, then click **Save** at the top of the page.

Alternative Architectures

As explained in the [Deployment Modes](#) section, you may want to deploy a second F5 device for various reasons. These alternative architectures require a few additional configuration steps.

Two F5 systems with ASAs deployed as a service pool

This solution is similar to the one explained in the section called, [Configuration for a Single F5 System with Cisco FirePOWER Services in L2 Mode](#). The only difference is that a second F5 device (the egress device) is introduced to offload re-encryption from the ingress device.

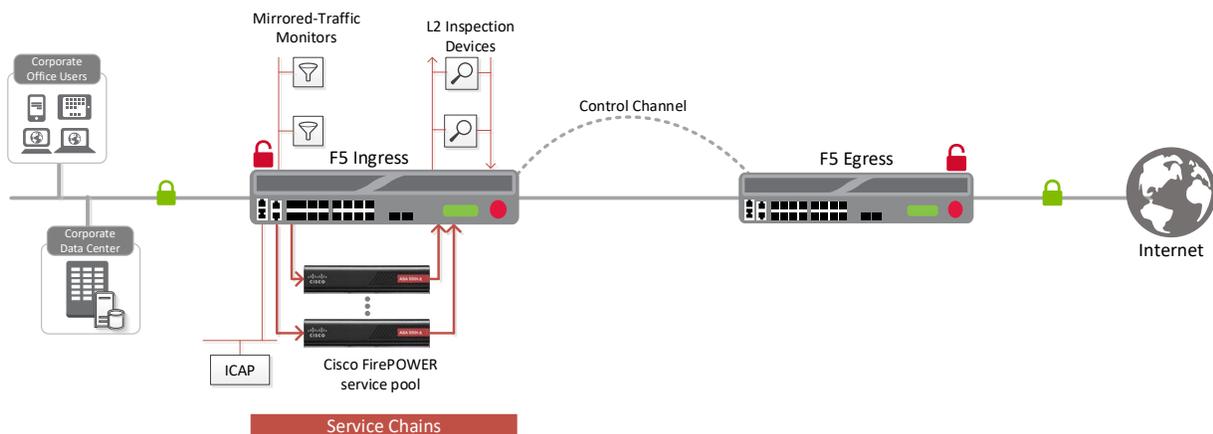


Figure 22: The SSL visibility solution with security service chaining using two F5 systems

Additional configuration steps

For this deployment scenario, you must configure SSL Orchestrator separately on the ingress and egress F5 devices to enable them to cooperate via a control channel. At least a /30 CIDR block is needed for IP connectivity of the control channel virtual servers on the F5 systems.

Configure the ingress F5 device: General properties

The ingress device is either an F5 device or a sync/failover device group where each client sends traffic. The ingress device is one or more ingress VLANs where the clients send traffic. The ingress device decrypts the traffic and then, based on protocol, source, and destination, classifies the traffic and passes each connection for inspection.

1. On the main tab, click **SSL Orchestrator > Configuration > General Properties**.
2. Enter information for the additional configurable fields that are specific to configuring separate ingress and egress devices. Follow the guidance below.

Question	User Input
Application Service Name	Enter a name without spaces or dashes for the SSL Orchestrator application service.
Do you want to setup separate ingress and egress devices with a cleartext zone between them?	Select Yes, configure separate ingress and egress BIG-IP devices .
Is this device the ingress or egress device?	Select This is the INGRESS device to which clients connect .
What is the EGRESS device Application Service name?	Enter the name of the SSL Orchestrator application service you intend to configure on the egress device. For the sake of ease, you can use the same SSL Orchestrator application service name on both the ingress and egress devices.
What is the IP address of the EGRESS device control-channel virtual server?	Enter the IP address of the control channel virtual server over on the egress device.
What IP address should THIS (ingress) device's control-channel virtual server use?	Enter the IP address of the virtual server for the control channel.
What is the control-channel pre-shared key?	Enter a pre-shared key (PSK) value to enable cryptographic protection of the service chain control channel between the ingress and egress F5 devices.

General Properties	
Application Service Name ?	SSLVisibility
Do you want to setup separate ingress and egress devices with a cleartext zone between them? ?	Yes, configure separate ingress and egress BIG-IP devices ▼
Is this device the ingress or egress device? ?	This is the INGRESS device to which clients connect ▼
What is the EGRESS device Application Service name? ?	SSLVisibility
What is the IP address of the EGRESS device control-channel virtual server? ?	192.168.10.2
What IP address should THIS (ingress) device's control-channel virtual server use? ?	192.168.10.1

Figure 23: Sample general configuration when ingress and egress devices will be separate

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

3. Continue configuration with the decryption zone configuration below.

Configure the ingress F5 device: Decryption zone to egress device configuration

1. Once you've completed the additional General Properties configuration, continue to scroll down the page and answer the following configuration question, using the guidance below.

Question	User Input
Are there parallel service devices in the decrypt zone?	Select No , send outbound traffic via the BIG-IP default route(s) . Or select Yes when firewalls will be sandwiched in the decrypt zone between the ingress and egress devices.

2. When done, click **Finished**, then click **Save** at the top of the page.

Configure the egress F5 device: General properties

The egress device is either an F5 device or a sync/failover device group that receives traffic that has traveled through the specified service chain and directs that traffic to the final destination. The ingress and egress devices also send each other control messages that can go through the decrypt zone or around it, if you configure a different path through the network. In either case, the messages are sent through TCP connections to port 245, at an IP address users specify, on each F5 system.

1. On the main tab, click **SSL Orchestrator > Configuration > General Properties**.
2. Enter information for the additional configurable fields that are specific to separate ingress and egress devices. Follow the guidance below.

Question	User Input
Application Service Name	Enter a name without spaces or dashes for the SSL Orchestrator application service.
Do you want to setup separate ingress and egress devices with a cleartext zone between them?	Select Yes , configure separate ingress and egress BIG-IP devices .
Is this device the ingress or egress device?	Select This is the EGRESS device which connects to a server .
What is the INGRESS device Application Service name?	Enter the name for the SSL Orchestrator application service configured on the ingress device.
What is the IP address of the INGRESS device control-channel virtual server?	Enter the IP address of the control channel virtual server on the ingress device.
What IP address should THIS (egress) device's control-channel virtual server use?	Enter the IP address of the virtual server for the control channel.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

What is the control-channel pre-shared key?	Enter a pre-shared key (PSK) value to enable cryptographic protection of the service chain control channel between the ingress and egress devices.
--	--

3. Continue configuration with the decryption zone configuration below.

Configure the egress F5 device: Egress device configuration

1. Once you've completed the additional General Properties configuration, continue to scroll down the page and answer this additional configuration question using the guidance below.

Question	User Input
Which VLAN(s) are part of the decrypt zone? (These bring traffic from the ingress device)	Select one or more VLANs where transparent-proxy ingress traffic will arrive.

Note: If you chose **Explicit proxy** in the General Properties section to answer the question “**Which proxy schemes do you want to implement?**” a separate explicit proxy configuration section displays here instead. In that case, choose the VLANs that explicit proxy will listen to, and enter the IP address and port number of the explicit proxy.

Configure the egress F5 device: Decrypt zone to ingress device configuration

1. Continue to scroll down the page and answer the question, **Are there parallel service devices in the decrypt zone?**
 - Select **No, send outbound traffic via the BIG-IP default route(s)**.
 - Or select **Yes** when firewalls will be sandwiched in the decrypt zone between the ingress and egress devices.
2. When done, click **Finished**, then click **Save** at the top of the page.

Two F5 systems with firewalls sandwiched in the decryption zone

In this case, the Cisco ASAs are deployed as a load balancing pool between the ingress and egress F5 systems in the decrypt zone and are *not* part of the service chains.

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

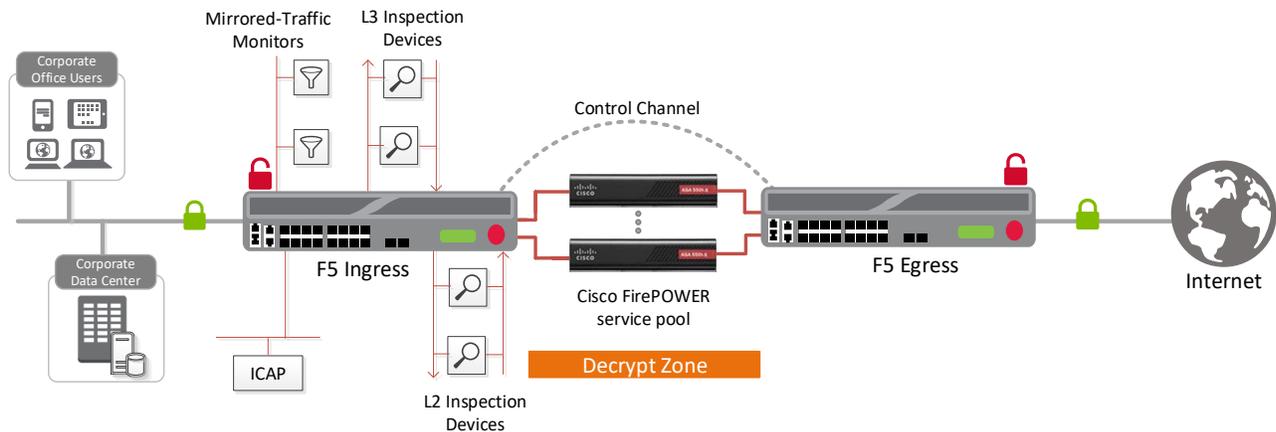


Figure 24: The SSL visibility solution with Cisco ASAs sandwiched between two F5 systems

Additional configuration steps

This scenario requires, as a prerequisite, that you have created a pair of VLANs—one on the ingress F5 device and another on the egress device for each Cisco ASA in the sandwiched pool. The IP addresses of each pair of VLANs will be from the same subnet if the firewalls are configured for L2 mode and from different subnets if the firewalls are configured in L3 mode.

Configure the ingress device: Decrypt zone to egress device configuration

1. On the main tab, click **SSL Orchestrator > Configuration > General Properties**. Scroll down to the additional configuration section specific to use of a decryption zone.

Decrypt Zone to Egress Device Configuration													
Are there parallel service devices in the decrypt zone? ?	<input type="checkbox"/> Yes, send outbound traffic via one or more service device(s) <input type="checkbox"/>												
What are the IPv4 decrypt zone gateway addresses? ?	<table border="1"> <thead> <tr> <th>Ratio</th> <th>IPv4 gateway address</th> <th>+</th> <th>-</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>192.168.20.1</td> <td>+</td> <td>-</td> </tr> <tr> <td>1</td> <td>192.168.30.1</td> <td>+</td> <td>-</td> </tr> </tbody> </table>	Ratio	IPv4 gateway address	+	-	1	192.168.20.1	+	-	1	192.168.30.1	+	-
Ratio	IPv4 gateway address	+	-										
1	192.168.20.1	+	-										
1	192.168.30.1	+	-										

Figure 25: Sample decrypt zone to egress device configuration

2. Answer the additional questions using the guidance below.

Question	User Input
Are there parallel service devices in the decrypt zone?	Select Yes, send outbound traffic via one or more service device(s) .

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

What are the IPv4 decrypt zone gateway addresses?	<p>Enter the IP address of the outward interface of the last layer 3 devices in the decrypt zone.</p> <ul style="list-style-type: none"> If the firewall is configured for L2 mode, the gateway IP address is the address on a VLAN on the next hop egress device. If the firewall is configured as an L3 hop, this IP address will be the address on a VLAN on the next hop firewall.
--	--

- Once you've entered an IP address, click **+** to add additional addresses. You can enter multiple gateways if you have multiple firewalls and want to load balance across them.
- Use the **Ratio** value to control the load balancing as desired.
- When done, click **Finished**, then click **Save** at the top of the page.

Configure the egress device: Decrypt zone to ingress device configuration

- On the main tab, click **SSL Orchestrator > Configuration > General Properties**. Scroll down to the additional configuration section specific to use of a decryption zone.

Decrypt Zone to Ingress Device Configuration			
Are there parallel service devices in the decrypt zone? ?	Yes, send outbound traffic via one or more service device(s) ▼		
What are the IPv4 decrypt zone gateway addresses? ?	Ratio	IPv4 gateway address	
	1	192.168.20.2	+ -
	1	192.168.30.2	+ -
What are the intranet networks (subnets)? ?	IPV4/IPV6 Subnet	IPV4/IPV6 Mask(CIDR)	
	10.10.10.0	255.255.255.0	+ -

Figure 26: Sample decrypt zone to ingress device configuration

- Answer the additional configuration questions using the guidance below.

Question	User Input
Are there parallel service devices in the decrypt zone?	Select Yes, send outbound traffic via one or more service device(s) .
What are the IPv4 decrypt zone gateway addresses?	Enter the IP address of the outward interface of the last layer 3 device in the decrypt zone.
What are the intranet networks (subnets)?	Enter the IP address and mask-length, in CIDR format, for intranet subnet masks. Typical IPv4 entries include 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16. Click + as needed to add additional addresses.

- When done, click **Finished**, then click **Save** at the top of the page.

Creating service chains to link services

Before you can set up service chains, you must configure all the services (inline, ICAP, or receive-only). By default, SSL Orchestrator steers traffic through all the security services. You can create a new service chain by defining the service list in the preferred order of services to which traffic should be steered.

Each service chain is linked to service chain classifier rules and processes specific connections based on those classifier rules, which look at protocol, source, and destination addresses. Additionally, service chains can include inline, ICAP, or receive-only services, as well as any decryption zones between separate ingress and egress devices.

1. On the main tab, click **SSL Orchestrator > Configuration > Policies**. The policies screen displays.
2. Under **Service Chains**, click **Add**.
3. Enter a **Name** for the service chain.
4. In the order you want SSL Orchestrator to use to steer the traffic, select the service **Type** (ICAP, inline or receive-only) and service **Name** and click **Add**. (See Figure 27.)
5. Repeat Step 4 until all services in the chain have been selected in the order you prefer.
6. When you're done with the service chain, click **Finished**.
7. Repeat Steps 2 through 6 to create multiple service chains.

Name		Services	
<input type="checkbox"/>	All	Type	Name
		receiveOnly	RSA
		icap	ICAP
		inlineService	Cisco
Show More			
<input type="checkbox"/>	PartnerNet	Type	Name
		inlineService	Cisco
		Inline Service	Cisco
			Add

Figure 27: Sample service chain configuration

Creating TCP service chain classifier rules

Before you create a TCP service chain classifier rule, you must [create one or more service chains](#). Service chain classifier rules then determine which service chains receive traffic. Each service chain classifier rule selects the specific chain to process ingress connections. Different classifier rules may send connections to the same chain. Each classifier has three filters that match the source IP address, the destination mode, and the application protocol. Filters can also overlap so that the classifier that matches best determines the service chain for a specific connection.

To avoid issues with privacy concerns and adhere to regulatory compliance, some organizations might need to enforce policies to bypass SSL destined to websites that expose personal user information, such as is the case for

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

banking, financial, or government sites. Classifier rules enable such policy implementation based on various context filters derived from a powerful classification engine. Finally, classifier rules can also be used to reject a connection if needed.

1. Once you've created a service chain, continue to scroll down the **Policies** page to **TCP Service Chain Classifiers**.
2. Click **Add** and create a classifier rule, making selections and completing each field using the guidance below.
3. The example below creates a sample TCP service chain classifier rule (as shown in Figure 28) to bypass SSL traffic originating from any internal client on 10.10.10.0 subnet in the corporate network and destined to any health care websites.

Configuration Field	User Input
Name	Enter a name for the TCP service classifier rule.
Phase	Select the SSL/TLS phase you want: No TLS: Match only non-TLS/SSL traffic. Pre-Handshake: Match TLS connections before any TLS handshake, which means you can allow a connection to bypass SSL inspection completely, without even trying to learn the real name of the remote server. Pre-handshake rules must reject or bypass any connections they match. TLS Handshake: Match only at the time of the TLS handshake and never match non-TLS traffic. The traffic is not checked again after the plaintext of a TLS connection becomes available. Normal: Match TLS connections at TLS handshake time and possibly again, more specifically, after SSL Orchestrator exposes the plaintext of the TLS connection (so you can manage HTTPS on non-standard ports, for example). Normal rules may also match non-TLS traffic (so, for example, a single rule can handle both HTTPS and HTTP traffic.) Select Normal in this sample SSL bypass configuration.
Protocol	Select the protocol to match: HTTP , MAIL , ALL , or Other . Select ALL to bypass all encrypted traffic.
Source	Select the source Type , either IP Address or Data Group , and then specify the filter Value . IP Address is either a traffic originating IP address or subnet. An explicit 0.0.0.0 will match all the traffic when IP address or subnet is not defined. Data Group is simply a user-defined group of related elements, such as a set of IP addresses. Refer to the AskF5.com resource on Data groups to learn more about data groups. Select IP Address as the source Type to match the connection originator and enter 10.10.10.0 in the Value field, then click Add .

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

<p>Destination</p>	<p>Select the destination Mode and specify the filter Type and Value, which may include:</p> <p>Address: Specify the traffic destination based on IP Address or Data Group (as with the source filter).</p> <p>Geolocation: Specify two-letter country and three-letter continent codes to match the destination IP against the local geolocation database.</p> <p>IPI: Specify the F5 IP Intelligence category or data group against which the destination IP address's reputation is validated. An IP Intelligence subscription is needed for the rule to evaluate against this database of known IP addresses with questionable reputations.</p> <p>Port: Specify the port or ports against which the destination port number should be matched. The value can be "any," one or more TCP port numbers, or ranges like 5557-5559 (use 0 or * to match all). The chief use of this mode is to control non-TLS traffic such as SNMP.</p> <p>URLF: Specify URL filtering (URLF) categories or a data group against which the destination URL will be matched. A URLF subscription is needed for the rule to evaluate against the URLF database.</p> <p>Name: Specify the domain name (with a unique name or using a wildcard) or data group against which the connection's hostname should be matched.</p> <p>DDB: Specify the DNS domain name (with a unique name or using a wildcard) against which the destination hostname indicated by the client in TLS Server Name Indication (SNI) is matched. Refer to RFC 6066 to understand the SNI extension for TLS. You may use DDB (dynamic domain bypass) to whitelist and bypass traffic to servers that cause TLS handshake problems or that require TLS mutual (client-certificate/smart-card) authentication. A URLF rule in the pre-handshake phase will match URL filtering categories associated with the TLS SNI hostname and otherwise behave like a DDB rule. See the example in Figure 28 below.</p> <p>Select URLF as the Destination Mode, Category as the Type, and Health and Medicine as the Value to match, if the connection is destined to any websites in the Health and Medicine category of the URLF database. Then click Add.</p>
<p>Service Chain</p>	<p>Select the name of a Service Chain (defined in the previous procedure) or an action—either Bypass or Reject.</p> <p>Select Bypass in the Service Chain selector to enforce a bypass action when both source and destination context filters match for an outbound connection.</p>

TCP Service Chain Classifiers ?

Add | Delete

Name	Phase	Protocol	Source	Destination	Service Chain
<input type="checkbox"/> Bypass	Normal	All	Type: IP Address, Value: 10.10.10.0	Mode: URLF, Type: Category, Value: Health and Medicine	Bypass

Finished Cancel

IP Address [] Add URLF [] Category [] Health and Medicine [] Add

Figure 28: Sample TCP service chain classifier

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

4. When your classifier rule configuration is complete, click **Finished**.
5. Repeat Steps 2 through 4 to create multiple TCP service chain classifiers.
6. If your answer to “**Do you want to pass UDP traffic through the transparent proxy unexamined?**” in the [General Properties](#) configuration was “**No, manage UDP traffic by classification,**” you will be presented with a **UDP Service Chain Classifiers** screen to create UDP rules similar to the TCP rules. Create and configure them following the same basic principles.
7. Finally, click **Deploy** at the top of the page to deploy the configured SSL Orchestrator.

Handling NAT

When a Cisco ASA with FirePOWER Services module is deployed as a service in the SSL Orchestrator service chain, it is no longer the Internet edge device. So performing the network address translation (NAT) on ASA is no longer advisable. It is also important to perform NAT of the client’s outbound traffic after it exits the ASA to the F5 egress for re-encryption. There are two ways to handle this:

- **Option A:** Implement NAT on the F5 system using the SNAT pool feature. (See Figure 29.) In this case, the NAT will be performed for the client’s outbound traffic on the egress of the F5 system. In the case of firewalls deployed as a sandwich pool using two F5 systems, NAT should be implemented on the egress F5 system.

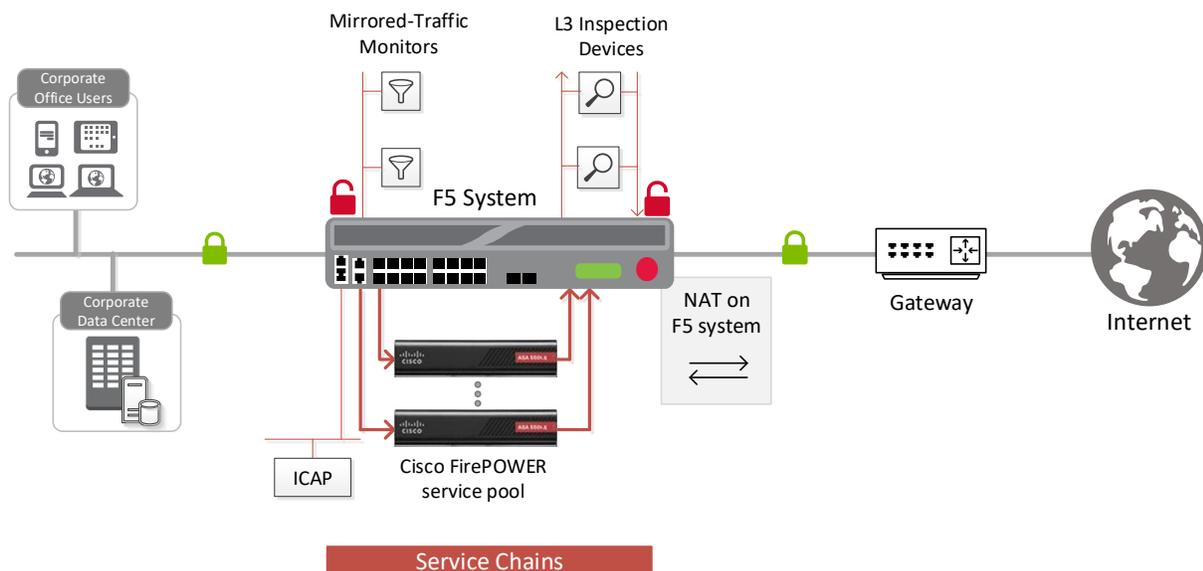


Figure 29: NAT on the F5 system (option A)

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

Traditionally, an ASA is often implemented on the perimeter to inspect/control access to multiple protocols, and not all of these protocols are supported by SSL Orchestrator. When this ASA is moved from the edge and configured in the service chain to inspect decrypted traffic, any unsupported protocol traffic that goes around SSL Orchestrator is not inspected and therefore potentially vulnerable.

The second option, Option B, represents the needed design change to overcome this challenge, as well as NAT recommendations.

- **Option B:** Segregate the ASA firewall capabilities and FirePOWER Services onto two different physical or virtual contexts, and implement NAT post re-encryption on the edge firewall while the inspection SFR module remains part of the F5 system in the service chain. (See Figure 30.) In this case, the F5 system can either hand off the re-encrypted packets to the edge firewall, or forward and re-route the traffic from the edge firewall to the gateway.

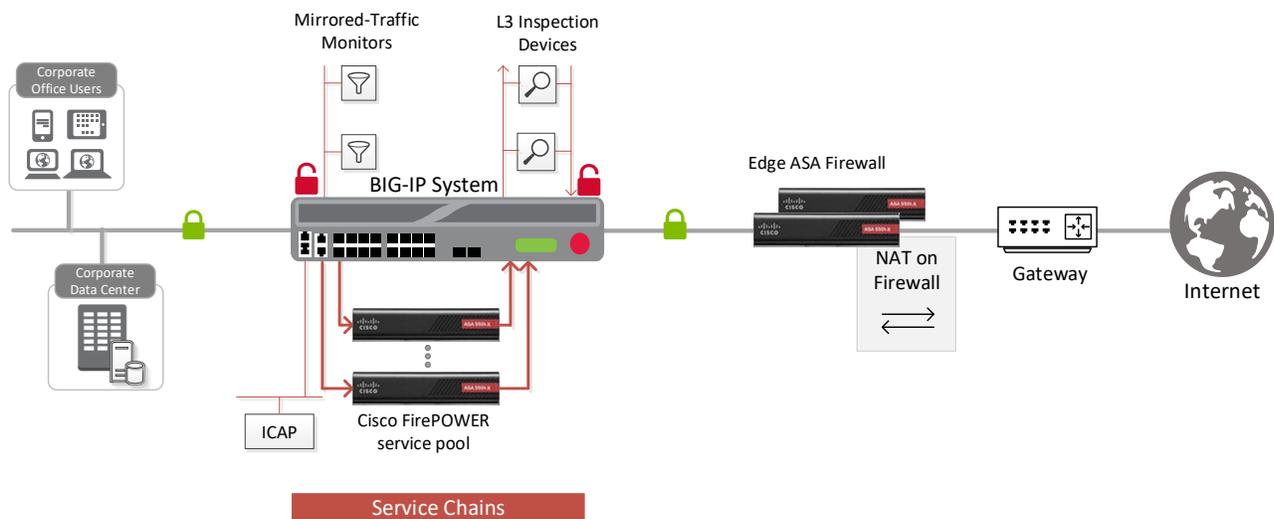


Figure 30: NAT on the Cisco ASA (option B)

Testing the Solution

You can test the deployed solution using the following options:

- **Server certificate test**

Open a browser on the client system and navigate to an HTTPS site, for example, <https://www.google.com>. Once the site opens in the browser, check the server certificate of the site and verify that it has been issued by the local CA set up on the F5 system. This confirms that the SSL forward proxy functionality enabled by SSL Orchestrator is working correctly.

- **Decrypted traffic analysis on the F5 system**

Perform a TCP dump on the F5 system to observe the decrypted cleartext traffic. This confirms SSL interception by the F5 device.

```
tcpdump -lnni eth<n> -Xs0
```

RECOMMENDED DEPLOYMENT PRACTICES

F5 and Cisco Firepower SSL Visibility with Service Chaining

- **Decrypted traffic analysis on the Cisco ASA**

From the web UI, go to **Monitoring > Packet Capture > Create**, and enable a **Packet Filter**. Create stages to capture packets, specify file names, and then click **OK**.

Download the captured file(s) and analyze the HTTP packets. The packet header and payload should be in clear text, indicating SSL decryption. It is very important to turn off packet capture once the job completes.

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 f5.com

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

Europe/Middle-East/Africa
emeainfo@f5.com

Japan
f5j-info@f5.com



©2018 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5.