F5® Herculon SSL Orchestrator™: Setup

Version 13.0

Table of Contents

What is F5 SSL Orchestrator?

What is F5 SSL Orchestrator

What is F5 SSL Orchestrator

F5® SSL Orchestrator™ provides an all-in-one appliance solution designed specifically to optimize the SSL infrastructure, provide security devices with visibility of SSL/TLS encrypted traffic, and maximize the efficient use of that existing security investment. This solution centralizes and consolidates SSL inspection across complex security architectures, allowing you flexible deployment options to decrypt and re-encrypt user traffic across the Internet and web-based applications. It supports policy-based management and steering of traffic flows to third-party security devices such as firewalls, intrusion prevention systems (IPS), anti-malware, data loss prevention (DLP), and forensics tools. It provides a wide range of SSL orchestration analytics that you can easily customize across multiple dimensions based on specified ranges of time.

The SSL Orchestrator single platform for unified inspection allows for the greatest flexibly without architectural changes to prevent new blind spots from emerging.

Some of the key functions include:

- Dynamic security service chaining that leverages context-based policies to efficiently deploy security, reduce administrative overhead, and effectively utilize security resources
- Centralized management of the SSL decrypt and re-encrypt function
- Inspection of all traffic for malware and data exfiltration with a multi-layered approach
- Flexible deployment modes to easily integrate the latest encryption technologies across your entire security infrastructure
- High availability with best-in-class load-balancing, health monitoring, and SSL offload capabilities

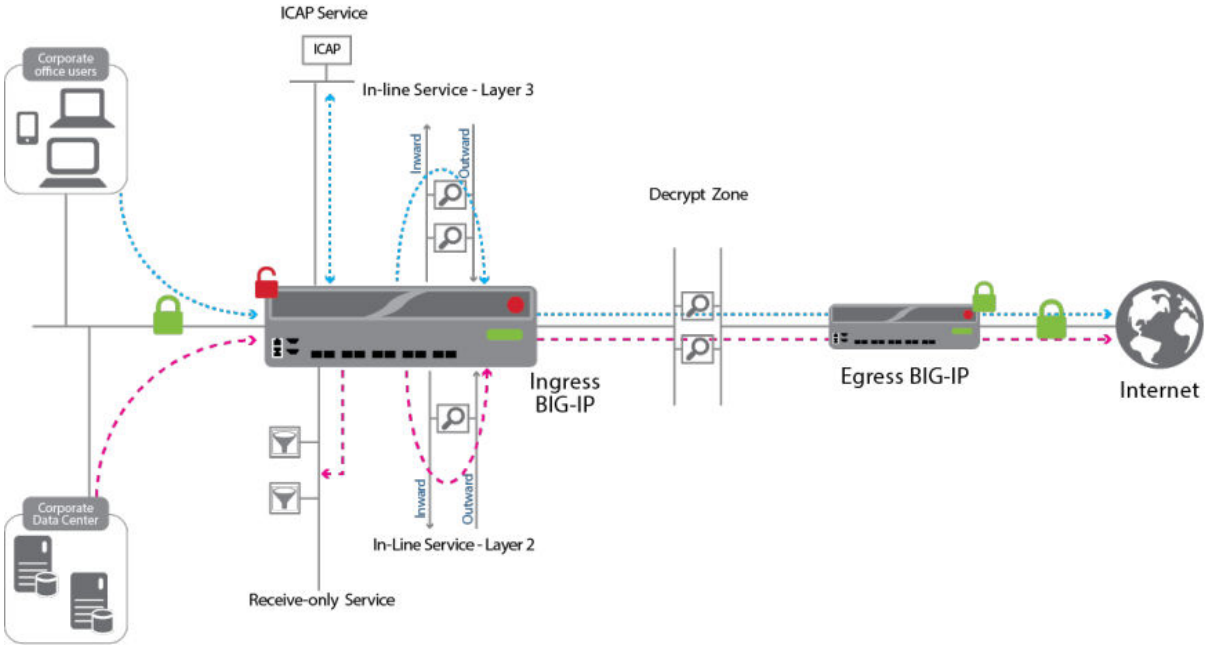

Figure 1: SSL Orchestrator solution

What is F5 SSL Orchestrator?

Terminology Used in this Implementation

Terminology Used in this Implementation

This section defines some of the terms used in this document.

- **Ingress device**

  The ingress BIG-IP® system is the device (or Sync-Failover device group) to which each client sends traffic. In the scenario where both ingress and egress traffic are handled by the same BIG-IP® system, ingress refers to the VLAN(s) where the client sends traffic. The ingress BIG-IP® system (or ingress VLAN(s)) decrypts the traffic and then based on protocol, source, destination, and so on, classifies it and passes each connection for inspection based on service chains you will configure (or allows certain connections to bypass service-chain processing based on your selections).
- **Egress device**

  The egress BIG-IP® system is the device (or Sync-Failover device group) that receives the traffic after a connection traverses the chosen service chain and then routes it to its final destination. In the scenario where both ingress and egress traffic are handled by the same BIG-IP® system, egress refers to the VLAN(s) where traffic leaves the BIG-IP® system to the Internet.
- **Inline services**

  Inline services pass traffic through one or more service (inspection) devices at Layer2 (MAC)/Bump-in-the-wire or Layer3 (IP). Each service device communicates with the ingress BIG-IP® device over two VLANs called *Inward* and *Outward*which carry traffic toward the intranet and the Internet respectively. You can configure up to ten inline services, each with multiple defined devices, using SSL Orchestrator.
- **Receive-only services**

  Receive-only services refer to services that only receive traffic for inspection, and do not send it back to the BIG-IP® system. Each receive-only service provides a packet-by-packet copy of the traffic (e.g. plaintext) passing through it to an inspection device. You can configure up to ten receive-only services using SSL Orchestrator. For more informationon receive-only services, see *Creating receive-only services for traffic inspection*.
- **ICAP services**

  Each ICAP service uses the ICAP protocol (*https://tools.ietf.org/html/rfc3507*) to refer HTTP traffic to one or more Content Adaptation device(s) for inspection and possible modification. You can add an ICAP service to any TCP service chain, but only HTTP traffic is sent to it, as we do not support ICAP for other protocols. You can configure up to ten ICAP services using SSL Orchestrator. For more information on ICAP services, see *Creating ICAP services*.
- **Service chains**

  SSL Orchestrator service chains process specific connections based on classifier rules which look at protocol, source and destination addresses, and so on. These service chains can include four types of services (Layer 2 inline services, Layer 3 inline services, receive-only services, and ICAP services) you define, as well as any decrypt zone between separate ingress and egress devices). For more information on service chains, see *Creating service chains to link services*.
- **Service chain classifier rules**

  Each service chain classifier rule chooses ingress connections to be processed by a service chain you configure (different classifier rules may send connections to the same chain). Each classifier rule has

four filters.The filters match source (client) IP address, destination (which can be IP address, IP Intelligence category, IP geolocation, domain name, domain URL Filtering category, or server port), and application protocol (based on port or protocol detection). Filters can overlap so the implementation chooses the classifier rule with the most specifc matches for each connection.

For more information on service chain classifier rules, see *Creating TCP service chain classifier rules* and/or *Creating UDP service chain classifier rules*.

- **Decrypt zone**

  A decrypt zone refers to the network region between separate ingress and egress BIG-IP® devices where cleartex data is available for inspection. Basically an extra inline service can be placed at the end of every service chain for additional inspection. You cannot configure a decrypt zone in the scenario where a single BIG-IP® system handles both ingress and egress traffic because the decrypt zone does not exist.

- **Transparent/Explicit Proxy**

  This implementation can operate in transparent and/or explicit proxy mode. A transparent proxy intercepts normal communication without requiring any special client configuration; clients are unaware of the proxy in the network. In this implementation, the transparent proxy scheme can intercept all types of TLS and TCP traffic. It can also process UDP and forward other types of IP traffic. The explicit proxy scheme supports only HTTP(S) per RFC2616. In addition, transparent proxy supports direct routing for policy-based routing (PBR) and Web Cache Communication Protocol (WCCP) that are dependent on networking services to support both protocols, while explicit proxy supports manual browser settings for proxy auto-config (PAC) and Web Proxy Autodiscovery Protocol (WPAD) that require additional iRule configurations (not included) to provide the PAC/WPAD script content.

- **Certificate Authority (CA) certificate**

  This implementation requires a Certificate Authority PKI (public key infrastructure) certificate and matching private key for SSL Forward Proxy. Your TLS clients must trust this CA certificate to sign server certificates.

- **SNAT**

  A SNAT (Secure Network Address Translation) is a feature that defines routable alias IP addresses that the BIG-IP® system substitutes for client IP source addresses when making connections to hosts on the external network. A **SNAT pool** is a pool of translation addresses that you can map to one or more original IP addresses. Translation addresses in a SNAT pool should not be self IP addresses.

- **Sync-Failover device group**

  A Sync-Failover device group (part of the Device Service Clustering (DSC®) functionality) contains BIG-IP® devices that synchronize their configuration data and failover to one another when a device becomes unavailable. In this configuration, a Sync-Failover device group supports a maximum of two devices.

Terminology Used in this Implementation

Configuring the System for F5 SSL Orchestrator

Overview: Configuring the system for F5 SSL Orchestrator

To set up your system for decrypting and encrypting outbound SSL/TLS traffic, you need to use the SSL Orchestrator Setup Wizard which initially guides you through basic minimal setup configuration. When you have completed the basic setup using the Setup Wizard, the SSL Orchestrator configuration utility will assist you with the rest of your configuration.

Configuring the System for F5 SSL Orchestrator

Using the SSL Orchestrator setup wizard

Before you start this task:

- Make sure you set up a management IP address, netmask, and default routing on your system.

Note: *If at any time during your configuration you need to return to the F5® SSL Orchestrator™ Setup Wizard, simply click the F5 logo in the upper-left corner of the Configuration utility, and on the Welcome screen, click the* ***Run the Setup Utility*** *link.*

The SSL Orchestrator Setup Wizard guides you through the basic, minimal setup configuration for F5®SSL Orchestrator™.

1. On the Welcome screen, click **Next**.
2. On the License screen, click **Activate**.
3. On the EULA screen, click **Accept**.
   The license activates and the system reboots for the configuration changes to take effect.
4. After the system reboots, click **Continue**.
5. On the Device Certificates screen, click **Next**.
6. On the Platform screen, for the **Management Port Configuration** setting, click **Manual**.

   The **Management Port** setting should include the management interface details that were previously created.
7. In the **Host Name** field, type the name of this system.

   The Host Name must be a fully qualified domain name.

   For example, `www.siterequest.com`.
8. In the User Administration area, type 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 system notifies you to log out and then log back in with your username and new password.

9. Click **OK**.
   The system reboots.
10. (Optional) On the Network Time Protocol (NTP) screen, in the **Address** field, type the IP address of the NTP server to synchronize the system clock with an NTP server, and click **Add**.
11. Click **Next**.
    The Domain Name Server (DNS) screen opens.
12. (Optional) To resolve host names on the system, set up the DNS and associated servers:
    a) For the **DNS Lookup Server List**, in the **Address** field, type the IP address of the DNS server and click **Add**.
    b) If you use BIND servers, add them in the **BIND Forwarder Server List**.
    c) For local domain lookups to resolve local host names, add them in the **DNS Search Domain List**.
    d) Click **Next**.

    The Internal VLAN screen opens.

    Note: *If you plan to later use the DNSSEC option in the iApp template, you must set up DNS using the SSL Orchestrator Setup Wizard. Otherwise, this step is optional.*

    ---

13. Specify the **Self IP** settings for the internal network:
    a) In the **Address** field, type a self IP address.
    b) In the **Netmask** field, type a network mask for the self IP address.
    c) For the **Port Lockdown** setting, retain the default value.
14. For the **VLAN Tag ID** setting, retain the recommended default value, **auto**.
15. For the **Interfaces** setting:
    a) From the **VLAN Interfaces** list, select an interface number.
    b) From the **Tagging** list, select **Tagged** or **Untagged**.
       Select **Tagged** when you want traffic for that interface to be tagged with a VLAN ID.
    c) Click **Add**.
16. Click **Next**.
    This completes the configuration of the internal self IP addresses and VLAN, and the External VLAN screen opens.
17. Specify the **Self IP** setting for the external network:
    a) In the **Address** field, type a self IP address.
    b) In the **Netmask** field, type a network mask for the self IP address.
    c) For the **Port Lockdown** setting, retain the default value.
18. In the **Default Gateway** field, type the IP address that you want to use as the default gateway to the external VLAN.
19. For the **VLAN Tag ID** setting, retain the recommended default value, **auto**.
20. Click **Next**.
    This completes the configuration of the external self IP addresses and VLAN.
21. On the Forward Proxy Certificate screen, do the following:
    a) In the **Certificate Name** field, select **Create New** and type a certificate name.
    b) In the **Certificate Source** field, select either **Upload File** and click **Choose File**, or select **Paste Text** and copy and paste your certificate source.
    c) In the **Key Source** field, select either **Upload File** and click **Choose File**, or select **Paste Text** and copy and paste your key source.
    d) From the **Security Type** list, select either **Normal** or **Password**.
22. Click **Next**.
23. On the Logging screen, under **Publisher Type**, select either **local** or **splunk**.

- If you select **local** as your **Publisher Type**, specify the **Destination** as either **local-db** or **local-syslog** and click **Next**.

  Note: *This determines the destination of your logs as being either a local database or a local syslog server.*

- If you select **splunk** as your **Publisher Type**:

a) For **Protocol**, select either **TCP** or **UDP**.
b) Type the **IP** address and the **Port** of the splunk server.
c) Click **Next**.

You are now ready to proceed to the second part of the configuration where you follow additional instructions to finalize your system for SSL Orchestrator. Refer to the *F5® Herculon SSL Orchestrator™: Setup* document for instructions.

Configuring the System for F5 SSL Orchestrator

Updating the SSL Orchestrator version

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

1. Open a web browser and go to *downloads.F5.com*
   You will need your credentials to login
2. Click **Find a Download**.
   The Select a Product Line screen opens.
3. In the **Herculon** F5 Product Family section, select **SSL Orchestrator**.
4. If necessary, select your BIG-IP® version from the list, and then click a product container from the list.
5. Accept the End User License agreement.
   The Select a Download screen opens.
6. From the list, select and download the latest version of the SSL Orchestrator zip file on to a location accessible from your system, and continue to follow the prompts to download the zip file.
7. Return to your SSL Orchestrator configuration utility.
8. Select **SSL Orchestrator** > **Updates**.
9. In the **File Name** field, click **Browse** and navigate to the file you saved onto your system and click **Open** to select it.
10. Click **Install**.
    The latest version of the SSL Orchestrator configuration utility is now installed. You system may reboot for change to take effect.

Configuring the System for F5 SSL Orchestrator

Backing up your BIG-IP configuration

For details, complete instructions, and other considerations for backing up and restoring the BIG-IP® configuration, see SOL 13132 on AskF5: *Backing up and restoring BIG-IP configuration files (11.x - 12.x)*

Before beginning the SSL Orchestrator configuration, or before you make substantial changes, we strongly recommend you back up the BIG-IP® configuration using the following steps. This allows you to restore the previous configuration in case of any issues.

1. On your system, click **System** > **Archives**.
2. To initiate the process of creating a new UCS archive (back up), click **Create**.
3. In the **File Name** box, type a name for the file. This name must be a unique name.
4. Click **Finished**.
5. To restore the configuration from a UCS archive, go to **System** > **Archives.**
6. Select the name of the UCS file you want to restore and click **Restore**.

Your BIG-IP® configuration is now safely restored.

Configuring the System for F5 SSL Orchestrator

Modifying your SSL Orchestrator configuration

We recommend that you backup your BIG-IP® configuration prior to making any changes to your SSL Orchestrator configuration. Refer to the "Backing up the BIG-IP Configuration" section of this document for more information.

You can modify your existing SSL Orchestrator configuration if you need to make changes.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**.
   A screen opens showing the network diagram and listing general properties.
2. Modify your configuration and then click **Deploy**.

Your existing configuration is now updated.

Configuring the System for F5 SSL Orchestrator

Deleting your SSL Orchestrator configuration

We recommend that you backup your BIG-IP® configuration prior to making any modifications to your SSL Orchestrator configuration. Refer to the "Backing up the BIG-IP configuration" section of this document for more information.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**.
   A screen opens showing the network diagram and listing general properties.
2. Click **Undeploy**.

Your entire configuration is now removed from your system.

Configuring the System for F5 SSL Orchestrator

Setting Up a Basic Configuration

Overview: Setting up a basic configuration

This section contains general information that the system needs before you can configure services and service chains. The SSL Orchestrator configuration utility will assist you with configuring logging settings, setting up ingress and egress devices as one system or separate systems, and configuring the system for transparent proxy and explicit proxy.

Setting Up a Basic Configuration

Configuring general properties

You must provide general information that the system needs so that you can then set up ingress and egress devices, create services and service chains, and create classifier rules using the F5 ® SSL Orchestrator™ configuration utility.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**.
   A screen opens showing the network diagram and listing general properties.
2. For the **Application Service Name** field, `ssloApp` is the default name for this configuration.
3. From the **Do you want to setup separate ingress and egress devices with a cleartext zone between them?** list, select one of the options:
   - If the same BIG-IP® system receives both ingress and egress traffic on different networks, use **No, use one BIG-IP device for ingress and egress**.
   - If you are configuring separate devices for ingress and egress traffic, use **Yes, configure separate ingress and egress BIG-IP devices**.
4. From the **Which IP address families do you want to support?** list, select **Support IPv4 only**.
5. From the **Which proxy schemes do you want to implement?** list, select whether the system operates in transparent proxy mode, explicit proxy mode, or both.
   - Use **Implement transparent proxy only** for the system to operate in transparent proxy mode. The transparent proxy scheme can intercept all types of TLS and TCP traffic. It also processes UDP traffic and forwards all other types of traffic. The transparent proxy requires no client configuration modifications.
   - Use **Implement both transparent and explicit proxies** for the system to operate in explicit and transparent proxy modes simultaneously.
   - Use **Implement explicit proxy only** for the system to operate in explicit proxy mode. The explicit proxy scheme supports only HTTP(S) per RFC2616. If you choose to configure an explicit proxy, assign a specific IP address and TCP port where the HTTP explicit-proxy clients connect.

6. From the **Do you want to pass UDP traffic through the transparent proxy unexamined?** list, select one of the options:
   - Use **Yes, pass all UDP traffic unexamined** to pass UDP traffic through without inspecting it.
   - Use **No, manage UDP traffic by classification** to configure specific service chain classifier rules for UDP traffic.

   This option is available only if you select **Implement transparent proxy only**.
7. From the **Do you want to pass non-TCP, non-UDP traffic through the transparent proxy?** list, select one of the options:
   - Use **Yes, pass non-TCP, non-UDP traffic** (such as IPsec, SCTP, OSPF, and so on) if you want the system to pass all traffic that is not TCP or UDP through the transparent proxy. If you choose this option, this traffic will not be classified or processed by any service chain.
   - Use **No, block all non-TCP, non-UDP traffic** (such as IPsec, SCTP, OSPF, and so on.) for the system to block all non-TCP and non-UDP traffic.

   This option is available only if you select **Implement transparent proxy only**.
8. From the **Which is the SSL Forward Proxy CA certificate?** list, select the Certificate Authority (CA) certificate that your clients will trust to authenticate intercepted TLS connections.
9. From the **Which is the SSL Forward Proxy CA private key?** list, select the corresponding private key.

   You imported the CA certificate and private key 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.
10. For the **What is the private-key passphrase (if any)?** field, type the private-key passphrase.

    If the key does not have a passphrase, leave the field empty.
11. From the **Which CA bundle is used to validate remote server certificates?** list, select the CA bundle that validates the remote server certificates.

    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.
12. From the **Should connections to servers with expired certificates be allowed?** list, select one of the two options to determine what happens with connections to servers with expired certificates:
    - Use **Yes, allow connections to servers with expired certificates** to allow connections to the servers that have expired certificates.
    - Use **No, forbid connections to servers with expired certificates** to prevent connections to servers that have expired certificates.

    Remote servers can present expired certificates. Allowing connections to servers with expired certificates can cause a security risk.
13. From the **Should connections to servers with untrusted certificates be allowed?** list, select one of the two options to determine what happens with connections to servers with untrusted certificates:
    - Use **Yes, allow connections to servers with untrusted certificates** to allow connections to the servers that have untrusted certificates.
    - Use **No, forbid connections to servers with untrusted certificates** to prevent connections to servers that have untrusted certificates.

    Remote servers can present untrusted certificates. Allowing connections to servers with untrusted certificates can cause a security risk.
14. If strict updates should protect the configuration, select the check box for **Should strict updates be enforced for this application?**.

    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. You should enable this setting to avoid misconfigurations that can cause an unusable application.

15. Click **Save**.

You have provided the basic configuration the system requires for SSL Orchestrator.

You can now set up ingress and egress devices, configure transparent or explicit proxies for the system, and create services, service chains, and classifier rules.

Setting Up a Basic Configuration

Configuring logging

Before configuring logging for SSL Orchestrator, complete all areas in General Properties.

You can generate log messages to help you monitor (and optionally debug) system activity. And you can choose the level of logging you want the system to perform. Log messages may be sent to one or more external log servers (preferred) and/or stored on the BIG-IP® device (less desirable because BIG-IP devices have limited log storage capacity).

1. On the main tab, click **SSL Orchestrator** > **Configuration**.
   The Logging Configuration section displays at the bottom of the screen.
2. From the **What SSL Intercept logging level do you want to enable?** list, select the level of logging the system performs.
   - Use **Errors. Log only functional errors** to log errors related to how SSL Orchestrator functions.
   - Use **Normal. Log connection data as well as errors** to log per-connection data in additional to functional errors.
   - Use **Debug. Log debug data as well as normal level data** to log debug data as well as connection data and functional errors. Because this logging level consumes more resources on the BIG-IP® system, use this mode only during setup or troubleshooting.
3. From the **Which Log Publisher will process the log messages?** list, select whether an existing log publisher object processes the log messages or does not process the log messages and sends the messages to syslog-ng. We strongly recommend that you use a Log Publisher for good system performance. The `syslog-ng` service is useful for Errors-only logging but is too slow for Normal or Debug logging when the system is used in production. A Log Publisher delivers log messages to one or more Log Destinations. Log Destinations may include Syslog, ArcSight, Splunk, and other log servers as well as the BIG-IP system's local log database. To use a Log Publisher, it must already be present on the system
   - Use **None (Send log messages to syslog-ng)** to send log messages to the system management plane syslog-ng subsystem. This option is not recommended for use in production systems.
   - Otherwise, from the list, select the Log Publisher you created. A Log Publisher delivers log messages to one or more Log Destinations. Log Destinations may include Syslog, ArcSight, Splunk, and other log servers.
4. From the **What kind of statistics do you want to record?** list, select the type of statistic the system records. This implementation can collect usage data for connections, service chains, services, and so on. The implementation can also record remote domain names and TLS cipher suites for TLS connections if you wish, but gathering such data consumes more system resources.

   Domain names are taken from remote server PKI certificates (or client SNI in the case of Dynamic Domain Bypass) and may include a wildcard. TLS cipher suites may not be recorded when a connection bypasses interception.

   If you choose to collect any statistics, the BIG-IP system starts saving extra data in memory for the use of integration with performance reporting systems like Splunk or BIG-IP® iStats integration.
   - Use **None** if you do not want the system to record statistics.
   - Use **Usage counters only (No remote-domain+cipher records)** to record usage counters only and not statistics on remote-domain and cipher records.

- Use **Usage counters and remote-domain+cipher records (may slow system)** to record both usage counters and remote-domain and cipher records. This option can slow performance on your system.

5. Click **Save**.

You have configured logging options and completed the basic F5® SSL Orchestrator™ configuration.

Setting Up a Basic Configuration

Configuring an ingress and egress device on one system

The ingress device is either a device or a Sync-Failover device group where each client sends traffic. The egress device is either a device or a Sync-Failover device group that receives traffic after a connection travels through the specified service chain and directs the traffic to the final destination.

If both the ingress and egress traffic are used by the same BIG-IP® system, 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.

If both the ingress and egress traffic are used by the same BIG-IP® system, the egress device is one or more egress VLANs where the clients receive traffic.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**.
   A screen opens showing the network diagram and listing general properties.
2. If you have only one BIG-IP® system, from the **Do you want to setup separate ingress and egress devices with a cleartext zone between them?** list, select **No, use one BIG-IP device for ingress and egress** .
3. From the **Which IP address families do you want to support?** list, select **Support IPv4 only**.
4. From the Ingress Device Configuration area, for the **Which VLAN(s) will bring client traffic to the transparent proxy?** setting, select one or more VLANs where transparent-proxy ingress traffic will arrive.
5. From the **How should a server TLS handshake failure be handled?** list, select whether you want the connection to fail or bypass the connection.
6. From the **DNS query resolution** list, select whether to permit the system to send DNS queries directly out to the Internet, or specify one or more local forwarding nameservers to process all DNS queries from SSL Orchestrator.
7. From the **Do you want to configure local/private DNS zones?** list, select whether you want to configure local or private DNS zones.
8. For the **Which local forwarding nameserver(s) will resolve DNS queries from this solution?** setting, type the IP address of local nameservers that will resolve all DNS queries from this implementation.
9. In the **List local/private Forward Zones** setting, type the IP address of one or more nameservers.
10. From the **Do you want to use DNSSEC to validate DNS information?** list, select whether you want to use DNSSEC to validate the DNS information.
11. In the Egress Device Configuration area, from the **Do you want to SNAT client IP addresses?** list, select whether you want to define SNAT addresses.
12. From the **Should traffic go to the Internet via specific gateways?** list, choose whether or not you want the system to let all SSL traffic use the default route, or if you want to specify Internet gateways (routers). If you chose to use specific gateways, you can also define the ratio of traffic sent to each device in the next step.

    - If you want outbound/Internet traffic out using the default route on the BIG-IP® system, select **No. Send outbound/Internet traffic via the default route**.

- If you want to define a list of gateways (routers) to handle outbound SSL traffic (and control the share of traffic each is given), use **Yes. Send outbound/Internet traffic via specific gateways**.

13. Click **Save**.

You have now configured an ingress device and an egress device located on one system.

This describes only the fields, lists, and areas needed to configure an ingress and egress device on one system. You should complete the other areas in General Properties before moving on to create services and service chains.

Setting Up a Basic Configuration

Configuring an ingress device (for separate ingress and egress devices)

The ingress device is either a 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**.
   A screen opens showing the network diagram and listing general properties.
2. From the **Do you want to setup separate ingress and egress devices with a cleartext zone between them?** list, select **Yes, configure separate ingress and egress BIG-IP devices**.
3. From the **Is this device the ingress or egress device?** list, select **This is the INGRESS device to which clients connect**.
4. For the **What is the EGRESS device Application Service name?** field, type the name of the device service.
5. For the **What is the IP address of the EGRESS device control-channel virtual server?** field, type the IP address of the service chain control channel virtual server over on the egress device.
6. For the **What IP address should THIS (ingress) device's control-channel virtual server use?** field, type the IP address of the virtual server for the service chain control channel on a VLAN.
7. For the **What is the control-channel pre-shared key?** field, type a pre-shared key (PSK) value to enable cryptographic protection of the service chain control channel between the ingress and egress devices.
8. From the **Which IP address families do you want to support?** list, select **Support IPv4 only**.
9. From the Ingress Device Configuration area, for the **Which VLAN(s) will bring client traffic to the transparent proxy?** setting, select one or more VLANs where transparent-proxy ingress traffic will arrive.
10. From the **How should a server TLS handshake failure be handled?** list, select whether you want the connection to fail or bypass the connection.
11. From the **DNS query resolution** list, select whether to permit the system to send DNS queries directly out to the Internet, or specify one or more local forwarding nameservers to process all DNS queries from SSL Orchestrator.
12. From the **Do you want to configure local/private DNS zones?** list, select whether you want to configure local or private DNS zones.
13. In the **Which local forwarding nameserver(s) will resolve DNS queries from this solution?** area, type the IP address of local nameservers that will resolve all DNS queries from this implementation.
14. In the **List local/private Forward Zones** area, type the IP address of one or more nameservers.
15. From the **Do you want to use DNSSEC to validate DNS information?** list, select whether you want to use DNSSEC to validate the DNS information.

16. In the Decrypt Zone to Egress Device Configuration area, for **Are there parallel service devices in the decrypt zone?**, select whether you want to send outbound traffic using the BIG-IP® system default route(s) or send outbound traffic through one or more service devices.
    - If the system will send the traffic through its default route to the internet, which must be configured to point to the egress BIG-IP® system, use **No, send outbound traffic via the BIG-IP default route(s)**.
    - If your configuration includes any Layer 3 systems in the decrypt zone that must receive the traffic, use **Yes, send outbound traffic via one or more service device(s)**.
17. From the **What are the IPv4 decrypt zone gateway addresses?**, type the IP addresses or the IPv4 decrypt zone gateway.

    Click the + button to add additional addresses.

    If you answered the previous question **Yes, send outbound traffic via one or more service device(s)**, you will type the IP address of the inward interface of the first Layer 3 device in the decrypt zone. You can enter multiple gateways if you have multiple systems and wish to load balance across them. If you do enter multiple addresses, you can also use the ratio value to control the load balancing. 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.
18. Click **Save**.

You have now configured an ingress device for a system configured for separate ingress and egress devices.

This describes only the fields, lists, and areas needed to configure an ingress device. You should complete the other areas in General Properties before moving on to create services and service chains.

Setting Up a Basic Configuration

Configuring an egress device (for separate ingress and egress devices)

The egress device is either a device or a Sync-Failover device group that receives traffic after a connection travels through the specified service chain and directs the traffic to the final destination. When users set up separate ingress and egress devices, they send each other control messages. These 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 BIG-IP® system.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**.
   A screen opens showing the network diagram and listing general properties.
2. From the **Do you want to setup separate ingress and egress devices with a cleartext zone between them?** list, select **Yes, configure separate ingress and egress BIG-IP devices**
3. From the **Is this device the ingress or egress device?** list, select **This is the EGRESS device to which connects to server**.
4. For the **What is the INGRESS device Application Service name?** field, type the name of the device service.
5. For the **What is the IP address of the INGRESS device control-channel virtual server?** field, type the IP address of the service chain control channel virtual server over on the egress device.
6. For the **What IP address should THIS (egress) device's control-channel virtual server use?** field, type the IP address of the virtual server for the service chain control channel on a VLAN.
7. For the **What is the control-channel pre-shared key?** field, type a pre-shared key (PSK) value to enable cryptographic protection of the service chain control channel between the ingress and egress devices.
8. From the **Which IP address families do you want to support?** list, select **Support IPv4 only**.

9. From the Egress Device Configuration area, in the **Which VLAN(s) are part of the decrypt zone? (These bring traffic from the ingress device)** setting, select one or more VLANs where transparent-proxy egress traffic will arrive.
10. From the **Do you want to SNAT client IP addresses?** list, select whether you want to define SNAT addresses.
11. From the **Do you want to use a SNAT Pool?** list, select whether you want to use a SNAT pool or SNAT auto map to translate addresses.
12. For **IPv4 SNAT addresses**, enter the SNAT addresses if you are using them.
13. From the **Should traffic go to the Internet via specific gateways?** list, select whether you want the system to let all SSL traffic use the default route, or if you want to specify Internet gateways (routers). If you chose to use specific gateways, you can also define the ratio of traffic sent to each device in the next step.
    - If you want outbound/Internet traffic out using the default route on the BIG-IP® system, use **No. Send outbound/Internet traffic via the default route**.
    - I,f you want to define a list of gateways (routers) to handle outbound SSL traffic (and control the share of traffic each is given) use **Yes. Send outbound/Internet traffic via specific gateways**.
14. For **What are the IPv4 outbound gateway addresses?**, type one or more IPv4 addresses of one or more exit gateways.

    Click the + button to add additional addresses.
15. In the Decrypt Zone to Ingress Device Configuration area, for **Are there parallel service devices in the decrypt zone?**, select whether you want to send outbound traffic using the BIG-IP® system default route(s) or send outbound traffic through one or more service devices.
    - If the system will send the traffic through its default route, which must be configured to point to the ingress BIG-IP® system, use **No, send outbound traffic via the BIG-IP default route(s)**.
    - If your configuration includes any Layer 3 systems in the decrypt zone that must receive the responses to traffic, use **Yes, send outbound traffic via one or more service device(s)**i.
16. For **What are the IPv4 decrypt zone gateway addresses?**, type the IP addresses or the IPv4 decrypt zone gateway.

    Click the + button to add additional addresses.

    If you answered the previous question **Yes, send outbound traffic via one or more service device(s)**, you need to enter the IP address of the outward interface of the last Layer 3 device in the decrypt zone. You can enter multiple gateways if you have multiple systems and want to load balance across them. If you do enter multiple addresses, you can also use the ratio value to control the load balancing. 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.
17. For **What are the intranet networks (subnets)?**, type the IP address and mask-length in CIDR format for intranet submasks.

    Click the + button to add additional addresses. Typical IPv4 entries include `10.0.0.0/8`, `172.16.0.0/12`, and `192.168.0.0/16`.
18. Click **Save**.

You have now configured an egress device for a system configured for separate ingress and egress devices.

This describes only the fields, lists, and areas needed to configure an egress device. You should complete the other areas in General Properties before moving on to create services and service chains.

Setting Up a Basic Configuration

Configuring the system for transparent proxy

You can configure F5® SSL Orchestrator™ to operate in transparent proxy mode only. A *transparent proxy* intercepts normal communication without requiring any special client configuration, so clients are unaware of the proxy in the network.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**.
   A screen opens showing the network diagram and listing general properties.
2. From the **Which IP address families do you want to support?** list, select **Support IPv4 only**.
3. From the **Which proxy schemes do you want to implement?** list, select **Implement transparent proxy only**.
4. From the **Do you want to pass UDP traffic through the transparent proxy unexamined?** list, select one of the options:
   - Use **Yes, pass all UDP traffic unexamined** to pass UDP traffic through without inspecting it.
   - Use **No, manage UDP traffic by classification** to configure specific service chain classifier rules for UDP traffic.
5. From the **Do you want to pass non-TCP, non-UDP traffic through the transparent proxy?** list, select one of the options:
   - Use **Yes, pass non-TCP, non-UDP traffic** (such as IPsec, SCTP, OSPF, and so on) if you want the system to pass all traffic that is not TCP or UDP through the transparent proxy. If you choose this option, this traffic will not be classified or processed by any service chain.
   - Use **No, block all non-TCP, non-UDP traffic** (such as IPsec, SCTP, OSPF, and so on) for the system to block all non-TCP and non-UDP traffic.
6. Click **Save**.

You have now configured SSL Orchestrator to work in transparent proxy mode.

This describes only the fields, lists, and areas needed to configure SSL Orchestrator to work in transparent proxy mode. You should also complete the other areas in General Properties before moving on to create services and service chains.

Setting Up a Basic Configuration

Configuring the system for explicit proxy

You can configure F5® SSL Orchestrator™ to operate in explicit proxy mode only. Explicit proxy in SSL Orchestrator requires manual configuration of the client and supports only HTTP(S) based on RFC2616.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**.
   A screen opens showing the network diagram and listing general properties.
2. From the **Which IP address families do you want to support?** list, select **Support IPv4 only**.
3. From the **Which proxy schemes do you want to implement?** list, select **Implement explicit proxy only**.
4. In the **On which VLAN(s) should the explicit proxy listen?** area, select one or more BIG-IP® VLANs where the explicit proxy listens.
5. For **What IPv4 address and port should the explicit proxy use?**, select the IPv4 address and port that the BIG-IP® system should use for the explicit proxy virtual server.
6. Click **Save**.

You have now configured SSL Orchestrator to work in explicit proxy mode.

This describes only the fields, lists, and areas needed to configure SSL Orchestrator to work in explicit proxy mode. You should also complete the other areas in General Properties before moving on to create services and service chains.

Setting Up a Basic Configuration

Creating Services, Service Chains, and Classifier Rules

Overview: Creating services, service chains, and classifier rules

This section decribes how to create inline services, ICAP services, receive-only services, service chains, and classifier rules.

Creating Services, Service Chains, and Classifier Rules

Creating inline services for service chains

Before you create inline services, complete the sections in the General Properties tab.

Inline services pass traffic through one or more service devices at Layer 2 or Layer 3. You use inline services in service chains, where each service device communicates with the BIG-IP® device, on the ingress side and over two VLANs. These VLANs route traffic toward the intranet and Internet, respectively.

Layer 3 inline services requires you to provide the IP address of the service devices from the present choices in the SSL Orchestrator™ configuration. If you are using Layer 3 inline services, this configuration sends and receives information from the services using a pre-defined set of addresses.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**, and then click the **Inline Services** tab. A screen opens showing the network diagram and the Inline Services settings.
2. For the **What is the IPv4 (CIDR/19) subnet-block base address?** field, type in the address block.

   For IPv4, F5 recommends the default block `198.19.0.0/19` to minimize the likelihood of address collisions.

   Note: *When using Layer 3 inline services, you must address your systems to match the required ranges. 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.*

   ---
3. Click **Add**.
4. In the **Name** field, type a name for your configuration.

   Use a short, unique name for this service. This name can contain 1 -15 alphanumeric or underscore characters, but must start with a letter. Letters are not case-sensitive.
5. From the **Service Type** list, select **Layer 2** or **Layer 3**.
6. In the **Interfaces** area, select the BIG-IP® system interface and VLAN tag for each VLAN pair.

Each Inward VLAN must be connected to the same Layer 2 virtual network from every device in the Sync-Failover Device Group, and each Outward VLAN likewise, but to a distinct Layer 2 virtual network.

If you choose to use the **Ratio** field, the BIG-IP® system distributes connections among pool members in a static rotation according to ratio weights that you define. In this case, the number of connections that each system receives over time is proportionate to the ratio weight you defined for each pool member or node. This number must be between 1-100.

For example, if you have five devices and you assign a ratio of **1** to the first three devices, and a ratio of **2** to the fourth device, and a ratio of **3** to the fifth device; the first three devices with a ratio of 1 each receive 1/8 of the traffic. The fourth device receives 1/4 of the traffic, and the fifth device receives 3/8 of the traffic.

7. Under Available Devices, from the **IP Address** field, select the IP address pairs of the Layer 3 devices.
8. From the **Translate Port for HTTP Traffic** list, select one of the options.
   - Use **No** if the connections should use their original destination ports.
   - Use **Yes to Port 80** to send all HTTP traffic through port 80.
   - Use **Yes to Port 8080** to send all HTTP traffic through port 8080.
   - Use **Yes to Port 8443** to send all HTTP traffic through port 8443.
9. From the **Connection Handling On Outage** list, select one of the following:
   - Use **Skip Service** to allow connections to skip the service you are configuring if all the devices in the service are unavailable.
   - Use **Reject Connection** for the system to reject every connection reaching the service when the service is down.
10. Click **Finished**.
11. Click **Save**.

Note:

Layer 3 devices need to follow a specific fixed addressing scheme. For each of the 10 possible layer 3 inline services, you need to use the following configuration (with **x** being 0-9 representing the inline service):

Inward interface:

- Address: `198.19.x.61` through `68` (for each of the load balanced Layer 3 devices)
- Netmask: `255.255.255.128`

Outward Interface:

- Address: `198.19.x.161` through `168` (for each of the load balanced Layer 3 devices)
- Netmask: `255.255.255.128`

Routes:

- Default Gateway: `198.19.x.245`
- Gateway to internal networks: `198.19.x.10` (unless SNAT is used)

While the base address can be changed if needed, F5 recommends leaving it set to the default: `198.19.0.0`.

You have now configured an inline service for SSL Orchestrator.

After creating more than one service, you can now create a service chain.

Creating Services, Service Chains, and Classifier Rules

Creating ICAP services

Before creating ICAP services, complete the sections in the General Properties tab.

ICAP services use the RFC3507 ICAP protocol to refer HTTP traffic to one or more content adaptation devices to inspect or modify. You can add an ICAP to any TCP service chain, but only HTTP traffic is sent to the chain. Additionally, you can configure up to ten ICAP services using the SSL Orchestrator configuration utility to load balance across them.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**, and then click the **ICAP Services** tab. A screen opens showing the network diagram and the ICAP Services settings.
2. Click **Add**.
3. In the **Name** field, type a name for your configuration.
4. For **ICAP Devices**, type an IP address and port number.
5. Click **Add**.
6. From the **Headers** list, select **Default** or **Custom** for the mode.

   To edit the headers, use **Custom**.
7. From the **TCP Connections** list, select F5®**OneConnect™** or **Separate**.

   Use **OneConnect** to reuse the TCP connections to ICAP servers, which processes multiple transactions. If your ICAP servers do not support multiple ICAP transactions per TCP connection, select **Separate**. OneConnect will then be disabled.
8. For **Request** and **Response**, type the ICAP request and response URI, defined by RFC3507, that are related to the ICAP server.
   For example,`icap://${SERVER_IP}:${SERVER_PORT}/REQMOD`.
9. For **Preview Max. Length (bytes)**, type the number of bytes that are in the maximum length for the ICAP preview.

   Bytes of content, up to the specified number, are sent to the ICAP server as a preview of each HTTP request or response. If you set the maximum preview length to zero (0), then requests and responses are streamed through the ICAP server. The largest value currently supported is 51200 (50KB).
10. From the **Server Failure Handling** list, select **Reset Connection** or **Next Service Chain**.
    - Use **Reset Connection** for the system to reset the connection to the client, discarding the request and response.
    - Use **Next Service Chain** for the system to let the request or response continue to the next service in the service chain.
11. From the **Send HTTP/1.0 Requests to ICAP** list, select how to send requests to the ICAP service.
    - Use `HTTP/1.0 & HTTP/1.1` to send both HTTP/1.0 and HTTP/1.1 requests to the ICAP service.
    - Use `HTTP/1.1 only` to send only HTTP/1.1 requests to the ICAP service. Any HTTP/1.0 requests are not inspected.
12. Click **Finished**.
13. Click **Save**.

You have now configured an ICAP service.

After creating more than one service, you can now create a service chain.

Creating Services, Service Chains, and Classifier Rules

Creating receive-only services for traffic inspection

Before you configure receive-only services, complete the sections in the General Properties tab.

Receive-only services only receive traffic for inspection and do not send the traffic back to the BIG-IP® system. Each receive-only service provides a packet-by-packet copy of the traffic passing through the service to an inspection device. You can configure up to ten receive-only services using the SSL Orchestrator configuration utility.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**, and then click the **Receive Only Services** tab.
   A screen opens showing the network diagram and the Receive Only Services settings.
2. Click **Add**.
3. In the **Name** field, type a name for your configuration.
4. In the **MAC Address** field, type the MAC address of the receive-only device.
5. In the **IP Address** field, type the nominal IP address for this device.

   Each receive-only device requires a nominal IP host address to identify the device in the BIG-IP® system.
6. From the **VLAN** list, select the VLAN where the receive-only device resides.
7. From the **Interface** list, select the associated BIG-IP® system interface.
8. Click **Finished**.
9. Click **Save**.

You have now created a receive-only service for SSL Orchestrator.

After creating more than one service, you can now create a service chain.

Creating Services, Service Chains, and Classifier Rules

Creating service chains to link services

Before you can set up service chains, you must configure multiple services such as inline, ICAP, or receive-only.

You can create service chains using previously-created services. A *service chain* is a list of services linked to service chain classifier rules. Service chains process specific connections based on classifier rules that look at protocol, source, and destination addresses. Additionally, service chains can include the following types of services, as well as any decrypt zones between separate ingress and egress devices:

- Layer 2 inline services
- Layer 3 inline services
- Receive-only services
- ICAP services

1. On the Main tab, click **SSL Orchestrator** > **Configuration**, and then click the **Policies** tab.
   A screen opens showing the network diagram and the Service Chain settings.
2. Click **Add**.
3. In the **Name** field, type a name for your service chain.

   Create a short name for this service chain. A service chain name may contain 1-15 alphanumeric or underscore characters and must start with a letter (not case-sensitive). Use spaces or commas to separate service names.

   Note: *You cannot use any of the keywords "all", "bypass", "reject", or "drop", nor the name of any (inspection) service you previously configured as a service chain name.*
4. In the **Services** area, select a **Type** and **Name**.
5. Click **Finished**.
6. Click **Save**.

You have now configured a service chain.

After you create a service chain, configure either TCP or UDP classifier rules.

Creating Services, Service Chains, and Classifier Rules

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 determine which service chains receive traffic. Each service chain classifier rule selects the specific chain to process ingress connections. Different classifier rules can send connections to the same chain. Each classifier has three filters that match the source IP address, the destination, and the application protocol. Filters can also overlap, so the best matching classifier determines the service chain for a specific connection. Finally, classifiers can reject a connection or allow it to bypass the service chain.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**, and then click the **Policies** tab.
   A screen opens showing a network diagram and the TCP Service Chain Classifiers settings.
2. In the TCP Service Chain Classifiers area, click **Add**.
3. In the **Name** field, type a name for this rule.
4. From the **Phase** list, select a phase for this classifier.
   - Use **Normal** if the rule may 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 nonstandard ports, for example). Normal rules may also match non-TLS traffic (so, for example, a single rule can handle both HTTPS and HTTP).
   - Use **No TLS** if the rules match only non-TLS traffic.
   - Use **Pre-Handshake** to have the rules match before any TLS handshake. This means the rules can allow a connection to bypass SSL inspection completely,without even trying to learn the real name of the remote server. All Dynamic Domain Bypass (DDB) rules must have **Phase** set to **Pre-Handshake**.
   - Use **TLS Handshake** rules to have the rules match only at TLS handshake time: they will never match non-TLS traffic, and they are not checked again after the plaintext of a TLS connection becomes available.
5. From the **Protocol** list, select the protocol of the connection based on the port number or protocol recognition.
6. In the **Source** area, select a **Type** and a **Value**.

   This option specifies the name of the Service Chain you configured that you want to use for this classifier rule. From the **Type** list, either or
   - Select **IP Address** and type the required IP address in the **Value** field.
   - Select **Data Group** and select the name of your data group from the **Value** list.
7. In the **Destination** setting, select a **Mode**, a **Type**, and a **Value**.

   This option specifies the destination of the connection. The value of this field is based on the selection you made for the mode.
   - From the **Mode** list, choose the mode you want to use for this classifier rule. The mode you choose determines the value you will use for the destination. You can choose one of the modes for each classifier rule:
     - For **Address**, the **Destination** filter you configure consists of one or more IP subnet or host addresses just like the **Source** filter.

- For **Geolocation**, the **Destination** you configure contains 2-letter country and 3-letter continent codes against which the IP Geolocation of the destination server is compared. The continent codes are: **CAF**=Africa, **CAN**=Antarctica, **CAS**=Asia, **CEU**=Europe, **CNA**=North America, **COC**=Oceania, **CSA**=South. The country codes are those of ISO 3166 alpha-2.
- For **IPI** (IP Inspection), the **Destination** you configure contains one or more IP Intelligence categories against which the destination IP address's reputation is matched. You must replace SPACE characters in names of IP Intelligence categories with underscores (_) before adding them to **Destination**.
- For **Port**, the **Protocol** value must be **All**. The **Destination** contains one or more TCP port numbers or ranges like 5557-5559 (use 0 or * to match all) against which the destination port number is matched. The main use of this mode is to control non-TLS traffic such as SSH.
- For **URLF** (URL Filtering), the **Destination** you configure is one or more URL Filtering categories against which the URL categorization of the destination server is compared. You must replace SPACE characters in names of URL Filtering categories with underscores (_) before adding them to **Destination**.
- For **DDB** (Dynamic Domain Bypass), the **Destination** you configure contains one or more DNS domain names (unique or wildcard) against which the destination hostname indicated by the client in TLS SNI is matched. This mode is special because it classifies traffic before the SSL Orchestrator implementation attempts any TLS handshake with the remote server (that is, in Match Phase Pre-handshake). You may use **DDB** to whitelist and bypass traffic to servers which cause TLS handshake problems or that require TLS mutual (client-certificate/smart-card) authentication. For **DDB**, the **Service Chain** value you select must be **Bypass** or **Reject**.

  For security, the DDB facility ensures the destination IP address for each bypassed connection corresponds to the allowed domain. DDB may replace the destination IP address supplied by the client with one freshly obtained from DNS.
- For **Name** (domain name), the **Destination** you will configure contains one or more DNS domain names (unique or wildcard) against which the connection's destination host name is matched.

- From the **Type** list, depending on which **Mode** you selected, choose either **IP Address**, **Data Group**, **Category**, or **Domain Name** (or there will be no selection required).
- From the **Value** list or field, depending on which **Type** and **Mode** you selected, choose a value from the list or type in the required information (hover your mouse over the field for tips on required information).

8. From the **Service Chain** list, select the name of the service chain you configured that you want to use for this classifier rule. This must be the name you gave a service chain or a special keyword:

   - **All** means a chain including all services: first receive-only services, then ICAP services, then in-line services.
   - **Bypass** lets the connection go to its destination without traversing any service chain.
   - **Reject** terminates the connection.

   By specifying service chain classifier rules, if more than one classifier matches a connection, the best-matching classifier determines the service chain for that connection (so the order of classifier rules in the list is not important). Classifiers can also reject a connection or let it bypass the service chain (bypass TLS interception). The default action applies to connections which do not match any classifier.

   This classifier is the element of the SSL Orchestrator implementation which selects the proper service chain to handle each connection. A *connection*" is a particular packet flow between client (source) and server (destination), identified by the 5-tuple of IP protocol (TCP or UDP), plus client (source) and server (destination) IP addresses and port numbers. The classifier has a set of rules for TCP connections, and another set of rules for UDP when UDP service chains are enabled. The classifier matches information describing each connection, such as its client and server IP addresses, against criteria specified in the classifier rules. For example, a classifier rule might match all connections

from clients homed on a certain IP subnet. Another classifier rule might match all connections going to servers in a certain country (using IP Geolocation).

9. Click **Finished**.
10. From the **What should happen to unmatched connections?** list, select how the system should handle unmatched connections.
11. Click **Save**.

You have now created a TCP service chain classifier rule.

Creating Services, Service Chains, and Classifier Rules

Creating UDP service chain classifier rules

Before you create a UDP service chain classifier rule, you must create one or more service chains.

Service chain classifier rules 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, and the application protocol. Filters can also overlap, so the best matching classifer determines the service chain for a specific connection. Finally, classifiers can reject a connection or allow it to bypass the service chain.

1. On the Main tab, click **SSL Orchestrator** > **Configuration**, and then click the **Policies** tab.
   A screen opens showing the network diagram and the UDP Service Chain Classifiers settings.
2. In the UDP Service Chain Classifiers area, click **Add**.
3. In the **Name** field, type a name for this rule.
4. From the **Protocol** list, select the protocol of the connection based on the port number or protocol recognition.
5. In the **Source** area, select a **Type** and type a **Value**.

   This option specifies the name of the Service Chain you configured that you want to use for this classifier rule.
6. In the **Destination** area, select a **Mode** and **Type**, and type a **Value**.

   This option specifies destination of the connection. The value of this field is based on the selection you made for the mode.
7. From the **Service Chain** list, select **All**.
8. Click **Finished**.
9. From the **What should happen to unmatched connections?** list, select how the system should handle unmatched connections.
10. Click **Save**.

You have created a UDP Service Chain Classifier rule.

Creating Services, Service Chains, and Classifier Rules

Setting up SSL Orchestrator in a High Availability Environment

Overview: Setting up SSL Orchestrator in a high availability environment

This section describes how to deploy SSL Orchestrator High Availability (HA). SSL Orchestrator HA configuration and deployment ensures a decrease in downtime and eliminates single points of failure. The deployment of SSL Orchestrator's HA works in concert with the BIG-IP® device groups support to sync the SSL Orchestrator specific configuration items and is transparent to the user. The deployment occurs after completing a configuration change and selecting Deploy. The deploy request is first routed to one of the devices in the HA device group. This first device configures the box where the request is received. After successful deployment on that box, the request is repeated on other BIG-IP devices. With SSL Orchestrator installed onto a dedicated system with failover, it automatically takes over in case of system failure. Data is synchronized between the two systems ensuring high availability and consistent protection.

Note: *SSL Orchestrator high availability deployment is only supported for use with the SSL Orchestrator iApp 2.1 version.*

Setting up SSL Orchestrator in a High Availability Environment

Assumptions and dependencies

- HA Setup: BIG-IP® HA (CMI) must be set to Active-Standby mode with network failover. See the BIG-IP Device Service Clustering: Administration document for detailed information on Active-Standby HA mode.
- HA Setup: If the deployed device group is not properly synced or .rpm packages are not properly syncing, make sure your HA self IP (for example, `ha_self`) Port Lockdown is not set to **Allow None**. On the Main tab, click **Network** > **Self IPs** and click on your `ha_self`. If Port Lockdown is set to **Allow Custom**, check that the HA network port 443 is open on self IP.
- BIG-IP HA Devices: Only manual sync is supported.
- BIG-IP HA Devices: For use with SSL Orchestrator iApp 2.1, the devices in each BIG-IP HA pair must be the same model and run the same version of TMOS (including any hotfixes). Except for the management interface, you must configure both devices to use the same arrangement of network interfaces, trunks, VLANs, self IPs (address and subnet mask), and routes. For example, if one BIG-IP is connected to a specific VLAN/subnet via interface 1.1, the other BIG-IP must also be connected to that VLAN/subnet via interface 1.1. If the BIG-IP configurations do not match, this solution will not deploy correctly and HA failover will not work.
- User Experience: False positive configuration conflicts in the HA environment must be ignored.
- User Experience: Deployment must be initiated from Active HA BIG-IP.
- User Experience: If a non HA environment is changed to an HA environment, the application must be redeployed. Similarly, if an HA environment is changed to a non HA environment, the application must be redeployed.

- User Experience: The SSL Configuration page (**SSL Orchestrator** > **Configuration**) for each peer device can be refreshed in order to see all modified changes.

Setting up SSL Orchestrator in a High Availability Environment

SSL Orchestrator high availability deployment

To ensure that your SSL Orchestrator HA deployment succeeds, it is critical that each deployment step, as well as the assumptions and dependencies, are closely followed for both boxes. In addition, adhere to all prerequisites, noting that if the systems in the device group are not configured consistently, the deployment synchronization process may fail.

- Installing an updated .rpm file
- Configuring the network for high availability
  - Configuring the ConfigSync and Failover IP address
  - Adding a device to the local trust domain
  - Creating a Sync-Failover device group
- Synchronizing the device group
- Setting up a basic configuration for deployment

It is also critical to test the deployment after configuration as some failures may not be reported in the UI.

Prerequisites

- You have made sure that the information used to configure your devices are identical on both boxes before configuring the network for high availability. Without identical information on both devices, noted throughout the following steps, the HA deployment may fail.
- You must have successfully setup a HA ConfigSync device group prior to starting configuration. Basic instructions are below. For more detailed instructions, refer to the *BIG-IP Device Service Clustering: Administration* document, section "Managing Configuration Synchronization".
- You have successfully installed the most current .rpm file on the first device (the Active device).
- You have already installed SSL Orchestrator with the appropriate license information using the SSL Orchestrator Setup Wizard (or the CLI) and made sure your device setup information is identical on both boxes:
  - While using the SSL Orchestrator Setup Wizard, you have noted the details used for NTP and DNS setup and made sure they will be identical on both boxes. You may verify duplication by selecting **System** > **Configuration** > **Device** > **NTP** (or **DNS**).
  - You have made sure that any certificates used in the configuration are copied to all devices.
  - You have made sure that information is identical on all devices. This information should include any of the following that are needed:
    - Client network
    - External network
    - Decrypt zone network
    - Decrypt zone control network
    - Networks providing access to ICAP devices and Receive-only devices.
  - You have made sure the log publishers are configured and named the same.
  - You have made sure all systems use the same interfaces for any services (if interface 1.1 is used to send traffic to an inline layer 2 device on system A, interface 1.1 must also be used on systems B, C, and D).

Note: *Do not attempt to duplicate the configuration by saving and restoring a user configuration set (UCS) file from one machine to the other or any other cloning approach. There are several IDs that are required to be unique that will also be duplicated, causing additional problems.*

Note: *For more detailed information on using the SSL Orchestrator Setup Wizard, see the "Using the SSL Orchestrator setup wizard" section.*

Setting up SSL Orchestrator in a High Availability Environment
Installing an updated .rpm file
Configuring the network for high availability
Synchronizing the device group
Setting up a basic configuration for deployment

Installing an updated .rpm file

Make sure you have the latest version of SSL Orchestrator. This will establish the version that will later appear on your other BIG-IP® HA peer device. After downloading the latest version of the SSL Orchestrator zip file from *downloads.F5.com*, return to your SSL Orchestrator configuration utility. See the section *Update the SSL Orchestrator version* for more detailed installation instructions.

1. Create a backup of your current configuration.
2. On the Main tab, click **SSL Orchestrator** > **Updates**.
3. In the **File Name** field, click **Browse** and navigate to the file you saved onto your system.
4. Click **Open** to select it.
5. Click **Install**.

Note: *Only install the iApp package (the *.rpg file) on the Active system. That system will copy it to the other systems in the ConfigSync group.*

Later, after a successful SSL Orchestrator HA deployment, you should verify that the same version appears on the BIG-IP HA peer device.

SSL Orchestrator high availability deployment

Configuring the network for high availability

On the Active device, specify the settings for VLAN HA and self IP addresses. If needed, configure all devices involved in the high availability group for HA.

Note: *This network will connect the various devices and must be a common layer-2 network between all devices.*

1. On the Main tab, click **Network** > **VLANs**. The VLAN List page appears.
2. Click **Create**. A new page appears to configure your new VLAN.
3. In the **Name** field of General Properties, enter the name (for example, `ha_vlan`).
4. For the **Interfaces** setting:
   a) From the **Interfaces** list, select an interface number.
   b) From the **Tagging** list, select **Tagged** for traffic, for that interface, to be tagged with a VLAN ID.
   c) Click **Add**. The interface number you selected will appear as a tagged service.
   d) Select **Finished**.
   e) Next to the F5 logo, your device status will appear showing **ONLINE (ACTIVE)** and **Standalone** with green indicators showing their status as up and running.

5. On the Main tab, click **Network** > **Self IPs** . The New Self IP Configuration page appears.
6. In the **Name** field, enter the self IP name (for example, `ha_self`).
7. In the **IP Address** field, enter the IP address for the device.
8. In the **Netmask** field, enter the netmask for the device.
9. In the **VLAN/Tunnel** field, select the VLAN name (`ha_vlan`).
10. Click **Finished**.

SSL Orchestrator high availability deployment
Configuring ConfigSync and failover IP addresses
Adding a device to local trust domain
Creating a sync-failover device group

Configuring ConfigSync and failover IP addresses

Before creating the device group, you should configure the configuration synchronization (ConfigSync) and Failover IP addresses for each BIG-IP® system in the device group. The ConfigSync address is the IP address that the system uses when synchronizing configuration with peer devices, and the Failover address is the IP address that the system uses for network failover.

1. On the Main tab, click **Device Management** > **Devices**. The Device List page appears with your current device showing in the list.
2. Click on your device in the device list. The device Properties page appears.
3. Select the **ConfigSync** tab. The ConfigSync Configuration section appears showing the Local Address of that device.
4. From the **Local Address** list, select the VLAN address (`ha_vlan`).
5. Click **Update**.
6. Select the **Failover Network** tab and click **Add**. The New Failover Unicast Address page opens. In the **Address** field, make sure that the VLAN address (`ha_vlan`) is present.
7. Click **Repeat**.
8. After the page refreshes, from the **Address** list, select the Management Address.

Note: *Connection Mirroring is not supported.*

---

9. Click **Finished**. The Failover Unicast Configuration section will list both the VLAN HA (`ha_valn`) and Management Address devices.

Configuring the network for high availability

Adding a device to local trust domain

Any BIG-IP® devices that you intend to add to a device group must first be members of the same local trust domain. When a BIG-IP device joins the local trust domain, it establishes a trust relationship with peer BIG-IP devices that are members of the same trust domain. For example, if you are creating a device group with two members, you must log in to one of the devices and join the other device to that system's local trust domain. The devices can then exchange their device properties and device connectivity information.

1. On the Main tab, click **Device Management** > **Device Trust**. The Local Domain page appears.
2. Select the **Device Trust Members** tab. The Peer and Subordinate Devices page appears.
3. Click **Add**. The **Device Trust** page appears with the **Retrieve Device Credentials (Step 1 of 3)** section.
4. In the **Device Type** field, select **Peer**.
5. In the **Device IP Address** field, enter the IP address of your device.

6. Click **Retrieve Device Information**. The **Verify Device Certificates (Step 2 of 3)** section appears.
7. Click **Device Certificate Matches**. The **Add Device (Step 3 of 3)** section appears.
8. In the **Name** field, enter the name of the device you are adding.
9. Click **Add Device**. Next to the F5 logo, the status of your device should show **ONLINE (ACTIVE)** and **Connected** with a green indicator next to it showing its active and connected status.

Configuring the network for high availability

Creating a sync-failover device group

This task establishes failover capability between two or more BIG-IP® devices. If an active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You perform this task on any one of the authority devices within the local trust domain.

1. On the Main tab, click **Device Management** > **Device Groups**. The New Device Group page appears.
2. Click **Create**.
3. In the **General Properties** section, do the following:
   a) In the **Name** field, enter the name of your device group.
   b) In the **Group Type** field, select **Sync-Failover** from the list.
4. In the **Configuration** section, do the following:
   a) In the **Members** field, select both available devices from the **Available** list and add them to the **Includes** list.
   b) In the **Sync Type** field, select **Manual with Incremental Sync**.

   Note: *You must do a manual sync. If you select* ***Automatic with Incremental Sync****, your HA deployment will fail.*
5. Click **Finished**.

The **Device Group List** page appears listing your new device group. The **ConfigSync Status** column will indicate `Awaiting Initial Sync`.

Configuring the network for high availability

Synchronizing the device group

This task synchronizes the BIG-IP® configuration data from the local device to the devices in the device group. This synchronization ensures that devices in the device group operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

1. Next to the F5 logo, click on **Awaiting Initial Sync**. The **Device Management Overview** page appears showing your Device Groups.
2. In the **Sync Issues** section, select **ha** to expand the **Devices** and **Sync Options** sections.
3. In the **Devices** section, make sure you select the device showing Changes Pending.
4. In the **Sync Options** section, select **Push the selected device configuration to the group**.
5. Click **Sync**.

You have now completed your SSL Orchestrator high availability deployment. Next, setup a basic configuration for deployment on your Active device.

SSL Orchestrator high availability deployment

Setting up a basic configuration for deployment

Refer to the "Setting Up a Basic Configuration" section for detailed instructions on completing the basic configuration on your Active device.

Note: *You must create identical information on each device before deploying the configuration.*

After deploying your configuration on the Active device, the configuration is automatically synchronized with all of the other devices in the device group. Since some errors may not be apparent, it is critical that you thoroughly test and diagnose the success or failure of the deployment. The following steps can be taken to test the system.

SSL Orchestrator high availability deployment

Task summary for diagnosing and fixing high availability deployment

Even though the potential for SSL Orchestrator HA deployment is low, thorough verification is recommended. If your HA deployment fails, attempt:

- Verifying deployment and viewing logs
- Verifying the .rpm file version on both devices
- Configuring general properties and redeploying
- Reviewing error logs and performing recovery steps

Setting up SSL Orchestrator in a High Availability Environment
Verifying deployment and viewing logs
Verifying the .rpm file version on both devices
Configuring general properties and redeploying
Reviewing error logs and performing recovery steps

Verifying deployment and viewing logs

Verify that all expected and required virtuals, profiles, and BIG-IP® LTM and network objects (route-domains, VLANs, self IPs) have been created on each device in the HA device group. These will be items beginning with the name given to the application (for example, if the application was named SSLO, verify that all of the items named | Summary SSL Orchestrator 13.0.0 | 9 SSLO_* are the same on all boxes). Ensure that the .rpm files are in sync, verify deployment with or without services, and review the following logs for failures:

- `/var/log/restnoded/restnoded.log`
- `/var/log/restjavad.0.log`

Note: *Because the initial device in the HA device group repeats the configuration requests and propagates the configuration to other BIG-IP devices, make sure you verify the initial configured device first, followed by each device in the HA device group. If the initial device deployment configuration fails, all other device configuration deployments will not successfully be configured.*

Task summary for diagnosing and fixing high availability deployment

Verifying the .rpm file version on both devices

After a successful SSL Orchestrator HA deployment, verify that the latest version of the SSL Orchestrator zip file is installed on both devices.

1. On the Main tab, click **SSL Orchestrator** > **Updates**.
2. Check the versions in the **Version** field.

If the versions are not identical, you must install an updated .rpm file and verify that both boxes are identically configured.

Task summary for diagnosing and fixing high availability deployment

Configuring general properties and redeploying

1. Remove all configurations present on all devices.

   Note: *You may want to restore a backup file instead, per device, to remove all current configurations.*

2. For all devices, individually configure each section in the iApp and select **Deploy**. Verify that all new objects are properly synced and deployed.

   Note: *If synchronization or deployment issues persist after deploying after each section, attempt to deploy after updating each item (instead of after each section) in the iApp and verify that all new objects are properly synced and deployed.*

   Note: *See the "Configuring general properties" section for more detailed information.*

Task summary for diagnosing and fixing high availability deployment

Reviewing error logs and performing recovery steps

1. Verify that all BIG-IP® LTM and network objects are present on each of the devices in the HA device group.
2. If the configuration deployment fails on each device, review the logs:
   - `/var/log/restnoded/restnoded.log`
   - `/var/log/restjavad.0.log`
3. Use the following REST GET command to determine the state of the deployed device block in the REST storage:
   - `curl -s -k -u admin:admin` https://localhost/mgmt/shared/iapp/blocks `| json-format`
4. Since failure scenarios can vary, after reviewing the logs, attempt the following recovery steps:
   a) Redeploy SSL Orchestrator.
   If this succeeds, you have recovered from the failure situation.
   b) Undeploy SSL Orchestrator.
   By undeploying, a cleanup of MCP objects on each of the boxes occurs while also cleaning up required data properties within the block stored in REST storage. If this succeeds, attempt to redeploy again.
   c) If redeploy or undeploy fails, do the following:
      1. From command line (back door), `run > touch /var/config/rest/iapps/enable`.
      2. Refresh the SSL Orchestrator menu UI.
      3. Select the deployed application from the list and delete the application.
      4. Redeploy and undeploy again.
      5. Once done, remove the file `rm -f /var/config/rest/iapps/enable`.
   d) If these recovery steps do not work, you may need to clean up the REST storage.

Note: *For more detailed information on setting up HA, see the BIG-IP Device Service Clustering: Administration document.*

Task summary for diagnosing and fixing high availability deployment

Using SSL Orchestrator Analytics

Overview: About SSL Orchestrator analytics

F5® SSL Orchestrator™ analytics provide a customizable view into your SSL orchestration statistics, and enable you to flexibly choose the information you want to view based on specified ranges of time that you can select and easily adjust. By leveraging the multiple options available, you can analyze dimensions individually, compare groups of dimensions and their statistics, and sort the charts and tables as you diagnose the performance and health of your system's SSL orchestration.

When you initially launch the analytics dashboard by clicking **SSL Orchestrator > Analytics > Statistics**, data that has been collected over the last hour is displayed in five line charts:

- Hit Count
- Client Bytes In
- Server Bytes In
- Client Bytes Out
- Server Bytes Out

This initial display of data is unfiltered and includes statistics generated for the following dimensions in tables:

- Client Cipher Names
- Client Cipher Versions
- Server Cipher Names
- Server Cipher Versions
- Virtual Servers
- Servers (the final destination)
- Actions

You can also use the SSL Orchestrator analytics Scheduled Reports to set up an automatic reporting schedule and later view any stored scheduled statistical records.

Using SSL Orchestrator Analytics

About analytics dashboard capabilities

The F5® SSL Orchestrator™ analytics dashboard has customizable options so that you can view statistics for your orchestration's top virtual servers, top sites bypassed, top sites decrypted, the most used client ciphers and protocols, and determine where the top server ciphers and protocols are used.

By customizing the timeline, charts, and tables, you can also:

- Adjust the timeline for the statistics displayed in the charts and tables to customize statics captured across a wide range of time, spanning between just a minute up to an entire year.
- Create comparison charts for two or more members of a dimension.
- Filter data across one or more dimensions.
- Sort table-based statistics for any dimension.

Using SSL Orchestrator Analytics

Timeline capabilities

The customizable timeline capabilities give you the ability to produce a statistical analysis based on a specified range of time. When you first open the analytics dashboard, the default refresh time is set at 5 minutes. You can change the refresh rate to several different settings by selecting from the **5 min.** list at the top of the screen.

- **1 min.**, **5 min.**, or **10 min.** options reset the refresh time in minutes.
- **Off** freezes the data refresh collection.

You can also update the statistics on demand by selecting **Refresh** above the timeline.

Using SSL Orchestrator Analytics

Customizing timeline capabilities

You can customize the range of time in which you would like to view data, in both the charts and tables. When you first launch the analytics dashboard, the default range of time is set at the **Last hour**. You can change the range of time that appears on the timeline to several different settings.

1. Click **SSL Orchestrator** > **Analytics** > **Statistics**.
2. Above the timeline, click the **Last hour** list.
3. Select the range of time you would like to view statistics across.
   - **Last hour**
   - **Last 4 hours**
   - **Last day**
   - **Last week**
   - **Last month**
   - **Last year**
   - **Year to date**
   - **All**

   The range of time you specified appears above the timeline bar.

4. Use the timeline sliders at each end of the timeline, to manually scroll back and forth, further narrowing in on ranges of time that are critical to the ongoing performance and health of your system's SSL orchestration.

   Once the sliders have been manipulated, you can also click and drag the bar above the timeline based on the newly specified timeline.

Using SSL Orchestrator Analytics

Chart capabilities

You can reorder the customizable line charts by dragging them up or down to a different place in the chart stack. You can also minimize them so to hide their information.

Using SSL Orchestrator Analytics

Customizing chart capabilities

Within each line chart, you can identify a specific day and time within the time range selected. You can also select a block of time to be further analyzed. By leveraging the multiple options available, you can analyze dimensions individually, compare groups of dimensions and their statistics, and sort the charts as you diagnose the performance and health of your system’s SSL orchestration.

1. Click **SSL Orchestrator** > **Analytics** > **Statistics**.
2. On the title bar of the line chart, place your cursor anywhere, click and hold to drag the chart up or down the chart stack.
3. On the right side of the title bar, select the minus sign (**-**) to minimize the chart.
4. Select the plus sign (+) to again maximize the chart.

Using SSL Orchestrator Analytics

Table capabilities

The customizable dimension tables can be reordered and expanded and minimized just like the line charts. By using the tab at the top of the table’s pane on the dashboard, you can expand the tables toward the center of the dashboard so that you can view all the table columns collecting statistics.

Using SSL Orchestrator Analytics

Customizing table capabilities

You can select each column within a table individually, and sort it to view the data in ascending or descending order.

1. Click **SSL Orchestrator** > **Analytics** > **Statistics**.
2. On the title bar of the table, place your cursor at the beginning of the table's title and click the three lines.
3. Select **Sort By**.
4. Choose one of the table columns to highlight:
   - Hit Count Per Second

- Client Bytes Out Per Second
- Duration
- Server Bytes In
- Server Bytes In Per Second
- Hit Count
- Server Bytes Out Per Second
- Client Bytes In
- Client Bytes In Per Second
- Client Bytes Out
- Server Bytes Out

You can also individually select each row within a table to update the statistics within each available chart. You can also launch a comparison chart based on the table and the column of data that you selected to sort by.

5. To highlight a table member, within the table, select that row.
6. Right-click the highlighted row and select **Sort By**.
7. Choose one of the table columns you would like to create a chart from, and right-click the highlighted row.
8. Select **Add Comparison Chart**.
   A comparison chart appears indicating the dimension, row member, and statistical column that you chose.

   Note: *You can add more comparison charts as needed, reorder new charts added to the chart stack, and minimize them as necessary.*

9. To remove the comparison chart, click the **X** in the far-right corner of the chart's title bar.
10. To reset the table without any member of the table highlighted, right-click the highlighted row within the table and select **Clear Selection**.

Using SSL Orchestrator Analytics

Charting bytes in, bytes out, and hit count over time

To see bytes in and out for your clients and servers, as well as viewing the hit count over a determined range of time.

1. On the Main tab, click **SSL Orchestrator** > **Analytics** > **Statistics**.

   The dashboard’s default line charts, showing data collected over the last hour, is displayed.
2. To decrease the time range that is being analyzed, adjust the sliders on the timeline at the top of the screen.
   The data found in the charts and tables will adjust based on the new time range specified.
3. To increase the time range that is being analyzed, click **Last Hour** and select a time range from the list.
4. To focus on specific members of a dimension, click the name of a dimension, on the right, to expand its details and select the members that interest you.
   For example, you might click **Servers** and select one or more servers from the table that displays.

Using SSL Orchestrator Analytics

Comparing statistics on the top virtual servers

Compare SSL Orchestrator statistics across virtual servers.

1. On the Main tab, click **SSL Orchestrator** > **Analytics** > **Statistics**.
   The screen displays charts on the left and dimensions, such as **Client Cipher Names**, **Virtual Servers**, and **Actions**, on the right.
2. On the right, expand the **Virtual Servers** dimension.
   By default, the virtual servers are sorted with the busiest at the top.
3. Select two or more virtual servers, then right-click and select **Add Comparison Chart**.
   A new chart displays above the default charts. By default, **Hit Count** is charted for the selected virtual servers.
4. To compare another statistic, in the chart legend click **Hit Count** and select a statistic.
   For example, select **Duration** or **Hit Count Per Second**.
   A line graph displays for the statistic you selected. The comparison chart remains available on your screen as long as you keep your browser open.
5. To remove the comparison chart from the screen, click the **X** in the upper right corner of the chart.

Using SSL Orchestrator Analytics

Viewing the top sites bypassed

You want to see the sites for which SSL Orchestrator has bypassed the most requests, performing no decryption.

1. On the Main tab, click **SSL Orchestrator** > **Analytics** > **Statistics**.
   The screen displays charts on the left and dimensions, such as **Client Cipher Names**, **Virtual Servers**, and **Actions**, on the right.
2. On the right, expand the **Actions** and **Servers** dimensions.
   Data for the dimensions display in tables. By default, those with the greatest hit count display at the top.
3. Rearrange the screen to display more columns of data in the tables:
   - To switch to a table-only view, click the icon.
   - To change the widths of the tables and the charts across the display, drag and drop the icon.
4. In **Actions**, select **Bypassed**.
5. In **Servers**, the servers on which SSL bypass has occurred the most frequently are at the top of the list.

Using SSL Orchestrator Analytics

Viewing the top sites decrypted

You can see the servers for which SSL Orchestrator decrypted the most requests.

1. On the Main tab, click **SSL Orchestrator** > **Analytics** > **Statistics**.
   The screen displays charts on the left and dimensions, such as **Client Cipher Names**, **Virtual Servers**, and **Actions**, on the right.
2. On the right, expand the **Actions** and **Servers** dimensions.
   Data for the dimensions display in tables. By default, those with the greatest hit count display at the top.
3. Rearrange the screen to display more columns of data in the tables:
   - To switch to a table-only view, click the icon.

- To change the widths of the tables and the charts across the display, drag and drop the  icon.

4. In **Actions**, select **Intercepted**.
5. In **Servers**, the servers on which decryption has occurred the most frequently are at the top of the table.

Using SSL Orchestrator Analytics

Viewing the most used client ciphers and protocols

You want to see which client ciphers and protocols are most used.

1. On the Main tab, click **SSL Orchestrator** > **Analytics** > **Statistics**.
   The screen displays charts on the left and dimensions, such as **Client Cipher Names**, **Virtual Servers**, and **Actions**, on the right.
2. On the right, expand **Client Cipher Names** and **Client Cipher Versions**.
   Data for the dimensions display in tables. By default, those with the greatest hit count display at the top.
3. To view data for the top cipher, from **Client Cipher Names**, select the top record.
   Charts update to display data for the selected cipher only. **Client Cipher Versions** displays only the protocols used by the selected cipher.
4. To view data for the top ten ciphers, from **Client Cipher Names** select the top ten records.

Using SSL Orchestrator Analytics

Finding where the top server ciphers and protocols are used

You want to see which server ciphers and protocols are most used and where they are used.

1. On the Main tab, click **SSL Orchestrator** > **Analytics** > **Statistics**.
   The screen displays charts on the left and dimensions, such as **Client Cipher Names**, **Virtual Servers**, and **Actions**, on the right.
2. On the right, expand the **Server Cipher Names** and **Server Cipher Versions** dimensions.
   Data for the dimensions display in tables. By default, those with the greatest hit count display at the top.
3. To view data for the top cipher, from **Server Cipher Names**, select the top record.
   Charts update to display data for the selected cipher only. **Server Cipher Versions** displays only the protocols used by the selected cipher.
4. To view data for the top ten ciphers, from **Server Cipher Names** select the top ten records.
5. To view the servers and virtual servers involved in the transactions, on the right expand **Servers** and **Virtual Servers**.
   The tables display only the servers and virtual servers where the cipher and protocol were used.

Using SSL Orchestrator Analytics

Scheduling reports

With the F5®SSL Orchestrator™ analytics Scheduled Reports feature you can set up an automatic reporting schedule and later view any stored scheduled statistical records.

To set up your SSL Orchestrator scheduled reports, click **SSL Orchestrator** > **SSL Orchestrator** > **Scheduled Reports** and click the + icon.

The Chart Schedule Properties screen opens where you can set up the reporting details and the frequency of automatically generated reports sent to a specified email address.

To view your scheduled reports, click **SSL Orchestrator > Analytics > Scheduled Reports**. The Scheduled Reports screen opens displaying any generated records and the following information:

- Module
- Name
- Send To
- Report
- SMTP Server
- Frequency
- Send Time
- Status

Using SSL Orchestrator Analytics

Legal Notices

Legal notices

Publication Date

This document was published on April 7, 2017.

Publication Number

MAN-0645-00

Copyright

Trademarks

For a current list of F5 trademarks and service marks, see *http://www.f5.com/about/guidelines-policies/trademarks*.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by one or more patents indicated at: *https://f5.com/about-us/policies/patents*.

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

Index

R

S

T

U

V