INTEGRATION GUIDE

BIG-IP SSL Orchestrator and McAfee Web Gateway

SSL/TLS Visibility for Advanced Threat Analysis and Prevention

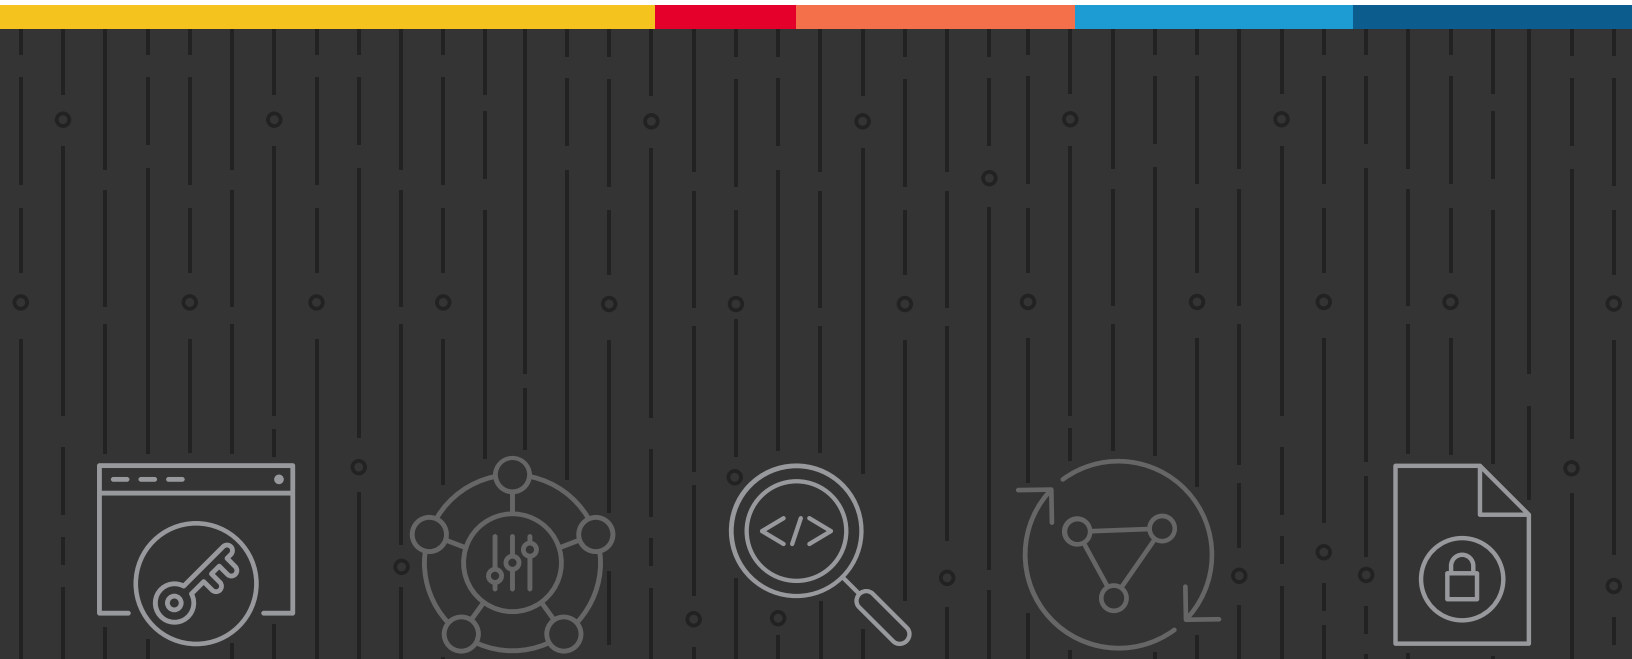

Table of Contents

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

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

An integrated F5 and McAfee solution solves these two SSL/TLS challenges. F5® BIG-IP® SSL Orchestrator® centralizes SSL/TLS inspection across complex security architectures, enabling flexible deployment options for decrypting and re-encrypting user traffic. It also provides intelligent traffic orchestration using dynamic service chaining and policy-based management. The decrypted traffic is then inspected by one or more McAfee Web Gateway (MWG) device, which can prevent previously hidden threats and block exploits. This solution eliminates the blind spots introduced by SSL/TLS and closes any opportunity for adversaries.

Note: BIG-IP SSL Orchestrator, with its ability to address HTTP proxy devices inside its decrypted inspection zone, allows MWG to provide optimal security functionality while offloading SSL/TLS and complex orchestration to the F5 system.

This guide provides an overview of the F5-McAfee joint solution and describes different deployment modes with reference to service chain architectures and recommended practices.

The F5 and McAfee Integrated Solution

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

BIG-IP SSL Orchestrator provides:

- **Multi-layered security.** To solve specific security challenges, security administrators are accustomed to manually chaining together multiple point products, creating a bare bone security stack consisting of multiple services. A typical stack may include components like data loss prevention (DLP) scanners, web application firewalls (WAFs),

intrusion prevention and detection systems (IPSs and IDSs), malware analysis tools, and more. In this model, all user sessions are provided the same level of security, as this "daisy chain" of services is hard-wired.

- **Dynamic service chaining.** Dynamic service chaining effectively breaks the daisy chain paradigm by processing specific connections based on context provided by the Security Policy, which then allows specific types of traffic to flow through arbitrary chains of services. These service chains can include five types of services: Layer 2 inline services, layer 3 inline services, receive-only services, ICAP services, and HTTP web proxy services.
  - **A service in BIG-IP SSL Orchestrator.** A service in BIG-IP SSL Orchestrator is defined as a pool of one or more same security devices. For example, a McAfee DLP ICAP service would include one or more McAfee DLP systems. BIG-IP SSL Orchestrator will automatically load balance the traffic to all the systems in a service.
  - **Health monitoring.** BIG-IP SSL Orchestrator provides various health monitors to check the health of the security devices in a service and handles failures instantly. For example, in a McAfee DLP ICAP service, should a system fail, BIG-IP SSL Orchestrator will shift the load to the active McAfee DLP systems. Should all the systems in the service fail, BIG-IP SSL Orchestrator will bypass the McAfee DLP ICAP service to maintain network continuity and maximize uptime.

- **Topologies.** Different environments call for different network implementations. While some can easily support SSL/TLS visibility at layer 3 (routed), others may require these devices to be inserted at layer 2. BIG-IP SSL Orchestrator can support all of these networking requirements with the following topology options:
  - Outbound transparent proxy
  - Outbound explicit proxy
  - Outbound layer 2
  - Inbound reverse proxy
  - Existing application
  - Inbound layer 2
- **Security policy.** The BIG-IP SSL Orchestrator security policy provides a rich set of context-aware methods to dynamically determine how best to optimize traffic flow through the security stack. Context can minimally come from the following:
  - Source and destination address/ subnet
  - URL filtering (URLF) and IP intelligence categories
  - Host and domain name
  - Destination port
  - IP geolocation
  - Protocol
- **Context engine for traffic classification.** BIG-IP SSL Orchestrator's context engine provides the ability to intelligently steer traffic based on policy decisions made using classification criteria, URL category, IP reputation, and flow information. In addition to

directing the traffic to service chains, customers can also use the context engine to bypass decryption to applications and websites, like financials, government services, health care, and any others, for legal or privacy purposes.

Figure 1: Context engine

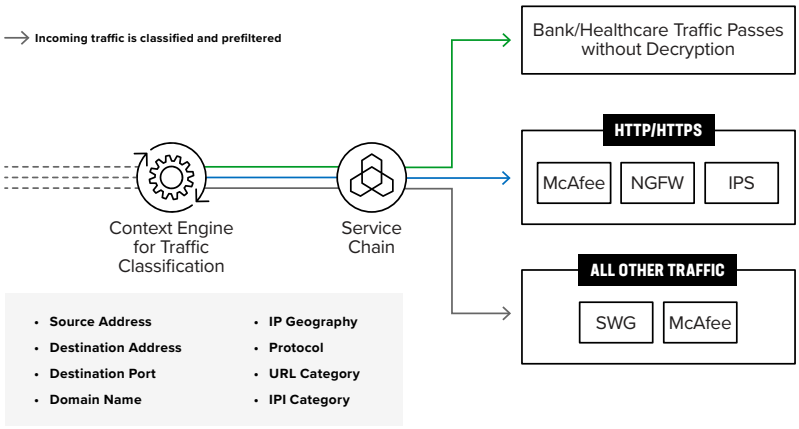

MWG provides:

- **Advanced anti-malware protection.** McAfee's anti-malware protection can identify known malicious files as well as analyze unknown files for hidden threats. Proactive intent analysis filters out previously unknown, or zero-day malicious content from web traffic in real time.
- **Intelligent sharing.** MWG creates and shares new file reputations for zero-day malware discovered by the gateway.
- **Application visibility and granular application control.** Application visibility grants full control over web applications such as those included in the Office365 suite, GSuite, Facebook, Webmail, Dropbox, etc.
- **Granular policy options.** McAfee's policy options can block individual web objects and file types based on any arbitrary Boolean combination of well over 200 transaction properties, including geolocation, category, reputation, user, user groups, and true file type.
- **Automated traffic analysis.** Automated analysis scans all web traffic in real time for both known and new malware, using dynamic reputation and behavior-based analysis on all web content. MWG also has the capability to quickly be adapted to new application features and associated security challenges like domain fronting and tenant restrictions.
- **Advanced threat analysis integration.** MWG integrates with McAfee Advanced Threat Defense, an advanced malware detection technology that combines customizable sandboxing with in-depth static code analysis.

SSL/TLS VISIBILITY: HOW DO WE DO IT?

F5's industry-leading full-proxy architecture enables BIG-IP SSL Orchestrator to install a decryption/clear-text zone between the client and web server, creating an aggregation (and, conversely, disaggregation) visibility point for security services. The F5® BIG-IP® system establishes two independent SSL/TLS connections—one with the client and the other with the server. When a client initiates an SSL/TLS connection to the server, the BIG-IP system intercepts and decrypts the client encrypted traffic and steers it to a pool of security devices for inspection before re-encrypting the same traffic to the server. The returned response from the server to the client is likewise intercepted and decrypted for inspection before being sent on to the client.

Figure 2: F5 full proxy architecture

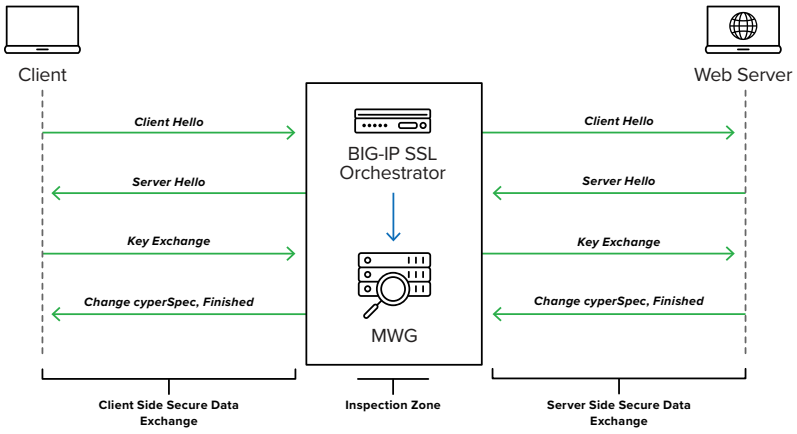

SSL/TLS ORCHESTRATION USING SECURITY SERVICE CHAINS

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

As shown in the figure below, BIG-IP SSL Orchestrator can load balance, monitor, and dynamically chain security services, including next-generation firewalls (NGFWs), DLP, IDS/IPS, WAF, and antivirus/malware, by matching the user-defined policies to determine whether to bypass or decrypt and whether to send to one set of security services or another. This policy-based traffic steering capability allows for better utilization of the existing security services investment and helps to reduce administrative costs.

Figure 3: Service chain

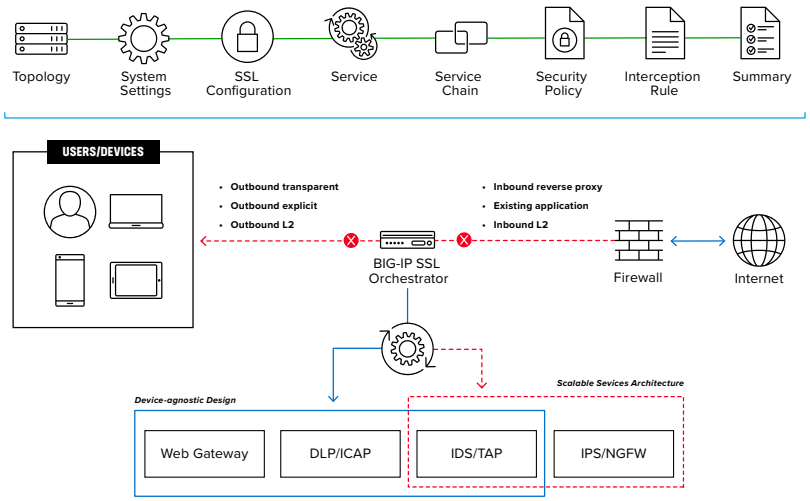

BIG-IP SSL Orchestrator enables you to apply different service chains based on context derived from a powerful classification engine. That context can come from:

- Source IP/subnet
- Destination IP/subnet
- IP intelligence category
- IP geolocation
- Host and domain name
- URLF category
- Destination port
- Protocol

Deployment Planning

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

SIZING

The main advantage of deploying BIG-IP SSL Orchestrator in the corporate security architecture is that the wire traffic now can be classified as “interesting” traffic, which needs to be decrypted by BIG-IP SSL Orchestrator for inspection by MWG, and “uninteresting” traffic is allowed to pass through or be processed differently according to other corporate policy requirements. This selective steering of only the interesting traffic to the firewall system conserves its valuable resources (as it need not inspect the entire wire traffic), maximizing performance.

As a result, it is important to consider the entire wire traffic volume to calculate the appropriate BIG-IP device size. The MWG system will require two interfaces on the BIG-IP systems (or one 802.1Q VLAN tagged interface) to allow traffic flow through logical inbound and outbound service interfaces.

Refer to the BIG-IP SSL Orchestrator data sheet and consider the following factors when sizing the BIG-IP system for the integrated solution:

- Port density.
- SSL/TLS bulk encryption throughput.
- System resources.
- The number of security services and devices in service chain.

Note: BIG-IP SSL Orchestrator has no specific port density requirement. Layer 3 must be layer 3 adjacent (routable), and layer 2 devices must be layer 2 adjacent (switched), and the BIG-IP system supports 802.1Q VLAN tagging, so a single interface can be logically divided into multiple VLANs. Security devices can connect to the BIG-IP system across a switched or routed architecture, so port density in this case is expandable. The only significant requirement is that inline security devices (layer 2, layer 3, and HTTP devices) must have separate physical or logical inbound and outbound interfaces.

LICENSE COMPONENTS

The BIG-IP SSL Orchestrator product line—the i2800, r2800, i4800, r4800, i5800, r5800, i10800, r10800, r10900, i11800, i15800, and High Performance Virtual Edition (HPVE) — supports this joint solution. The F5® VIPRION® platform and F5® VELOS® platform are also supported. BIG-IP SSL Orchestrator devices ship with an installed base module that provides both SSL/TLS interception and service chaining capabilities. Please contact your local F5 representative to further understand the licensing and deployment options.

Unless otherwise noted, references to BIG-IP SSL Orchestrator and the BIG-IP system in this document (and some user interfaces) apply equally regardless of the F5 hardware or virtual edition (VE) used. The solution architecture and configuration are identical.

Optionally, customers can add the functionality of:

- An **F5 URLF subscription** to access the URL category database.
- An **F5® IP Intelligence Services subscription** for IP reputation service.
- **A network hardware security module (HSM)** to safeguard and manage digital keys for strong authentication.
- **F5® Secure Web Gateway Services** to filter and control outbound web traffic using a URL database.
- **F5® BIG-IP® Access Policy Manager® (APM)** to authenticate and manage user access.
- **F5® BIG-IP® Advanced Firewall Manager™ (AFM)** to protect against denial-of-service.
- **F5® BIG-IP® Advanced WAF®** to protect against common vulnerabilities (CVEs) and web exploits, targeted attacks, and advanced threats.

- **F5® BIG-IP® Local Traffic Manager™ (LTM)** add-on software license mode. This solution is supported on all F5® BIG-IP® iSeries® and older F5 hardware platforms and has no specific restrictions on additional F5 software modules (including the above software services). This option is suited for environments that need to deploy BIG-IP SSL Orchestrator on an existing BIG-IP device or have other functions that must run on the same device.

TRAFFIC EXEMPTIONS FOR SSL/TLS INSPECTION

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

- Guest VLANs.
- Applications that use pinned certificates.
- Trusted software update sources.
- Trusted backup solutions.
- Any lateral encrypted traffic to internal services to be exempted.

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

Traffic exemptions based on URL category might include bypasses (and thus no decryption) for traffic from known sources of these types of traffic, including (but not limited to):

- Financial
- Health care
- Government services

CERTIFICATE REQUIREMENTS

Depending on the direction of flow there are different certificate requirements.

Outbound traffic flow (internal client to Internet)

An SSL/TLS certificate and associated private key—preferably a subordinate certificate authority (CA)—on the BIG-IP system are needed to issue certificates to the end host for client-requested external resources that are being intercepted. To ensure that clients on the corporate network do not encounter certificate errors when accessing SSL/TLS-enabled websites from their browsers, this issuing certificate must be locally trusted in the client environment.

Inbound traffic flow (Internet client to internal applications)

Inbound SSL/TLS orchestration is similar to traditional reverse web proxy SSL/TLS handling. It minimally requires a server certificate and associated private key that matches the host name external users are trying to access. This may be a single instance certificate, or wildcard or subject alternative name (SAN) certificate if inbound BIG-IP SSL Orchestrator is defined as a gateway service.

ARCHITECTURE RECOMMENDED PRACTICES

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

- Deploy inline. Any SSL/TLS visibility solution must be in-line to the traffic flow to decrypt perfect forward secrecy (PFS) cipher suites such as elliptic curve Diffie-Hellman encryption (ECDHE).
- Deploy the BIG-IP systems in a sync/failover device group, which includes an active/standby pair with a floating IP address for high availability (HA).
- Every MWG in the service pool must be dual homed on the “to-service” (BIG-IP to MWG) and “from-service” (MWG back to BIG-IP) VLANs with each BIG-IP system in the device sync/failover device group. This can be physically separate interfaces or a single 802.1Q tagged VLAN for logical separation.
- Further interface redundancy can be achieved using the Link Aggregation Control Protocol (LACP). LACP manages the connected physical interfaces as a single virtual interface (aggregate group) and detects any interface failures within the group.
- Unlike with some competing solutions, the BIG-IP systems do not need physical connections to the MWG. The BIG-IP system requires only layer 3 reachability to the MWG. In slow networks, however, we recommend deploying the services not more than one hop away.

SECURITY RECOMMENDED PRACTICES

SSL/TLS orchestration generally presents a new paradigm in the typical network architecture. Before, client-server traffic passed encrypted to inline security services, which then had to perform their own decryption if they needed to inspect that traffic. Integrated with BIG-IP SSL Orchestrator, now ALL traffic to a security device is decrypted—including usernames, passwords, social security and credit card numbers, etc. It is therefore highly recommended that security services be isolated within a private, protected enclave defined by BIG-IP SSL Orchestrator. It is technically possible, however, to configure BIG-IP SSL Orchestrator to send the decrypted traffic anywhere that it can route to, but this is a dangerous practice that should be avoided.

IP ADDRESSING

The recommended approach to integrating security devices is to physically move them to an isolated enclave of BIG-IP SSL Orchestrator. This generally requires re-addressing inline layer 3 security services. The following assumes this approach is being taken, and the example IP addresses represent a local/internal addressing scheme. As previously stated, it is entirely possible to configure BIG-IP SSL Orchestrator to send decrypted traffic anywhere that it can route to, but this is a dangerous practice.

When MWG is deployed as either a transparent proxy or explicit proxy, F5 recommends configuring its IP addresses for connected to-service and from-service interfaces from private addressing subnets provided by BIG-IP SSL Orchestrator. The default subnets are derived from an RFC2544 CIDR block of 198.19.0.0, which improves security and minimizes the likelihood of address collisions.

For example, you can configure an MWG to use the IP address 198.19.96.10/25 on its to-service interface and 198.19.96.130/25 on its from-service interface. The table below explains the IP addresses that you need to configure when deploying multiple MWGs in a service pool.

Devices	To-Service IP	From-Service IP	Gateway
MWG 1	198.19.96.10/25	198.19.96.130/25	198.19.96.245
MWG 2	1198.19.96.11/25	198.19.96.131/25	198.19.96.245
MWG N	198.19.96.x/25	198.19.96.x/25	198.19.96.245

The MWG would then default route back to BIG-IP SSL Orchestrator on a defined address in the from-service subnet. In the example above, that IP address is 198.19.96.245.

Note: A /25 network (255.255.255.128) subdivides a regular /24 network into two subnets and is a simple way to maximize local protected IP addressing for security services. These are the IP subnets and schemes that BIG-IP SSL Orchestrator uses by default, but any addressing scheme can be used:

198.19.96.1 - 198.19.96.126

198.19.96.129 - 198.19.96.254

Additionally, while it can, BIG-IP SSL Orchestrator does not by default source NAT (SNAT) the client source address across inline layer 3 services. Plus, it is usually favorable for inline security devices to see the true client IP address. Inline layer 3 services may therefore also need a static route back to the to-service side of BIG-IP SSL Orchestrator for IPs/subnets that match the client-side network. For example:

Client-side network:	10.20.0.0/24
BIG-IP SSL Orchestrator service to-service self:	198.19.96.7/25
BIG-IP SSL Orchestrator service from-service self:	198.19.96.245/25

In this example, the client would be coming from the 10.20.0.0/24 subnet, an IP subnet foreign to the MWG. It may be necessary to create a static route that directs MWG to forward any return traffic destined to this subnet back to BIG-IP SSL Orchestrator to-service self IP address (ex. 198.19.96.7).

Initial Setup

Initial setup includes configuration of MWG and setup of BIG-IP SSL Orchestrator. Once these steps are complete, you can proceed to configuration for the specific deployment scenario you choose.

CONFIGURE MWG PREREQUISITES

Before the MWG can receive traffic from BIG-IP SSL Orchestrator, there are a few basic configurations that must be completed. Any and all licenses should be applied, and basic system setup should be completed. Along with many other settings, the system setup will include configuration of the hostname and Domain Name Systems (DNS). The system hostname will be configured, as well as the IP address, subnet mask, and hostname, for the management interface. Additional interfaces will be configured further on in this guide.

CONFIGURE BIG-IP SSL ORCHESTRATOR PREREQUISITES

Before you begin configuring BIG-IP SSL Orchestrator, there are a few prerequisite operations that need to be addressed.

Define client side and outbound (e.g., Internet) side VLANs and self-IPs

For BIG-IP SSL Orchestrator in a layer 3 (routed or explicit proxy) topology, the BIG-IP system must be configured with appropriate client-facing and outbound-facing VLANs and self-IPs. The VLANs define the connected interfaces, and the self-IPs define the respective IPv4 and/or IPv6 subnets. In the BIG-IP system, under the Network menu, configure the client side and outbound side VLANs and self-IPs appropriately.

Import CA certificate and private key

For BIG-IP SSL Orchestrator in an outbound traffic topology, a local CA certificate and private key are required to resign the remote server certificates for local (internal) clients. In the BIG-IP system, under System > Certificate Management, import the required CA certificate and private key. Ensure that internal clients trust this local CA.

Update the BIG-IP SSL Orchestrator application

Periodic updates are available for BIG-IP SSL Orchestrator. (If you are upgrading from a previous major version, refer to the BIG-IP SSL Orchestrator setup guide for the recovery procedure.)

To download the latest update:

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

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

Configuring BIG-IP SSL Orchestrator Integration with MWG

BIG-IPSSL Orchestrator is configured to send decrypted traffic to an inline MWG. BIG-IP SSL Orchestrator handles both decryption and re-encryption of HTTPS traffic, with an inspection zone installed between the ingress and egress. Decrypted traffic is steered to a service pool of MWG devices. You can also deploy the BIG-IP system as a device sync/failover device group (including an HA pair) with a floating IP address for HA.

The MWG can be configured as either a transparent proxy or explicit proxy inside the inspection zone.

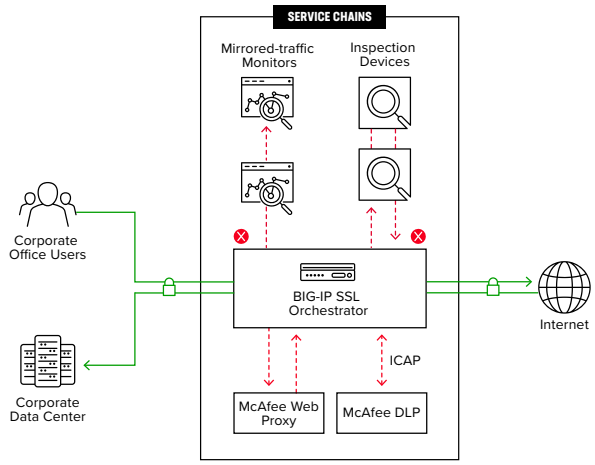

Figure 4: Sample traffic flow

How traffic flows in this deployment:

- Client traffic arriving at the ingress side of the BIG-IP system is classified, and interesting HTTPS traffic is decrypted as part of the SSL/TLS handling process.
- BIG-IP SSL Orchestrator steers the decrypted traffic through the load balanced MWG service pool as part of a service chain of potentially multiple types of security services.
- The HTTP traffic is inspected by the MWG services for any hidden threats before sending that traffic back to the BIG-IP system.

- The BIG-IP system orchestrates the decrypted traffic through other services in the chain before it aggregates and re-encrypts the traffic, which is then routed to the next destination.

Note: Inside the BIG-IP SSL Orchestrator inspection zone, all traffic passing to MWG is unencrypted HTTP. It is a security recommended practice to protect this device and the network between it and the BIG-IP SSL Orchestrator device. It is therefore recommended that the MWG devices be moved to the secure enclave created by BIG-IP SSL Orchestrator. This brings with it a few architectural changes.

If the MWG was previously installed in the network to service client traffic, as either a transparent or explicit proxy, BIG-IP SSL Orchestrator must now fulfill that role. If the MWG was configured as an explicit proxy, BIG-IP SSL Orchestrator must then be configured as an explicit proxy and clients must communicate with BIG-IP SSL Orchestrator as their explicit proxy gateway. If the MWG was configured as a transparent proxy, BIG-IP SSL Orchestrator must then be configured as a transparent proxy and routed client traffic must now pass through it. MWG devices inside the inspection zone can be explicit or transparent, irrespective of the BIG-IP SSL Orchestrator proxy mode.

If the MWG device was previously handling explicit proxy user authentication, BIG-IP SSL Orchestrator must now fulfill this role. BIG-IP APM can be provisioned to provide the required explicit forward proxy authentication functionality.

CONFIGURE MWG

The following are the minimum requirements to configure an MWG device for integration with BIG-IP SSL Orchestrator. Please refer to McAfee documentation for additional product-specific information.

Configure interfaces

The MWG device must be physically or logically two-armed. In other words, it must have separate physical or logical to-service and from-service interfaces. This can either be two separate physical interfaces or a single 802.1Q VLAN tagged interface (a single interface with two tagged VLANs).

To complete the interface configuration in the MWG GUI, navigate to **Appliances > (this appliance) > Network Interfaces**. Relevant settings are described below:

- **Hostname:** Enter a unique hostname for this MWG appliance.
- **Default Gateway (IPv4):** Traffic will route from the MWG appliance back to BIG-IP SSL Orchestrator. The IP address used here will be the “from-service” F5 BIG-IP self-IP (ex. 198.19.96.245).

Under the “Enable these network interfaces” section:

- Enable the MWG “to-service” interface. This is the interface that receives decrypted HTTP traffic from BIG-IP SSL Orchestrator. Configure its IPv4 settings manually and supply an IP address in the pre-defined subnet (ex. 198.19.96.10). Enter the appropriate subnet mask (ex. 255.255.255.128).
- Enable the MWG “from-service” interface. This is the interface that sends decrypted HTTP traffic from MWG back to BIG-IP SSL Orchestrator. Configure its IPv4 settings manually and supply an IP address in the pre-defined subnet (ex. 198.19.96.131). Enter the appropriate subnet mask (ex. 255.255.255.128).

Note: BIG-IP SSL Orchestrator supports handling of both IPv4 and IPv6 traffic flows. To allow IPv6 traffic to flow through the MWG, add IPv6 addresses in the MWG configuration, then create an IPv6 version of the MWG service in BIG-IP SSL Orchestrator.

Optionally configure 802.1Q VLANs

If two separate data plane interfaces are not available, it is also possible to configure a single interface with two 802.1Q tagged VLANs.

From the MWG UI, under **Appliances > (this appliance) > Network Interfaces**, select and enable the raw interface to use.

- Configure the IPv4 and IPv6 settings for this interface as Disabled.
- Click the "Add VLAN..." button and supply a unique VLAN ID in the dialog box.
- Click the "Add VLAN..." button again and supply a unique VLAN ID in the dialog box.
- Click to enable the first VLAN and configure its IPv4 settings manually. Supply an IP address in the pre-defined subnet (ex. 198.19.96.10). Enter the appropriate subnet mask (ex. 255.255.255.128).
- Click to enable the second VLAN and configure its IPv4 settings manually. Supply an IP address in the pre-defined subnet (ex. 198.19.96.131). Enter the appropriate subnet mask (ex. 255.255.255.128).

Configure routes

The MWG devices will require two routes:

1. **A gateway route to send user traffic outbound:** Traffic will route from the MWG appliance back to BIG-IP SSL Orchestrator. The IP address used here will be the "from-service" F5 BIG-IP self-IP (ex. 198.19.96.245). This is defined in the MWG UI under **Appliances > (this appliance) > Network Interfaces**, in the "Default gateway (IPv4)" setting.
2. **A static return route:** BIG-IP SSL Orchestrator does not SNAT traffic across inline security devices by default, so the source address passing through the MWG will be foreign and will need a static return route to define the path back to the BIG-IP system on the inbound side of the MWG. For example:

Client-side network:	10.20.0.0/24
BIG-IP SSL Orchestrator service to-service self:	198.19.96.7/25
BIG-IP SSL Orchestrator service from-service self:	198.19.96.245/25

In this example, the client would be coming from the 10.20.0.0/24 subnet, an IP subnet foreign to the MWG. It is, therefore, necessary to create a static route that directs MWG to forward any return traffic destined to the subnet back to the BIG-IP SSL Orchestrator service inbound self IP address (ex. 198.19.96.7). This is defined in the MWG UI under **Appliances > (this appliance) > Static Routes**.

Configure DNS

DNS settings are minimally required to allow the MWG devices to talk to remote services, including DNS, license, and engine updates. There are generally two options for DNS. It can either pass through the management interface or the data plane outbound interface. DNS is configured in the MWG UI under **Appliances > (this appliance) > Domain Name Services**.

Configure the web proxy service

The services attached to BIG-IP SSL Orchestrator are opaque to the external environment. They are protected, isolated, and do not interact outside of the internal connectivity with the BIG-IP system. Most important, services (devices connected to the BIG-IP system) and topologies (how BIG-IP SSL Orchestrator consumes network traffic) are independent of one another. For example, MWG can be deployed as an explicit or transparent proxy inside the BIG-IP SSL Orchestrator inspection zone, while the deployed topology can be explicit forward proxy, transparent forward proxy, or even in a layer 2 bump-in-the-wire mode. With this in mind, the simplest, most efficient, and most flexible configuration option for the MWG in the BIG-IP SSL Orchestrator inspection zone is as an explicit proxy, irrespective of the BIG-IP SSL Orchestrator topologies defined.

The following settings will detail how to configure MWG as an explicit proxy. Please refer to the appropriate MWG documentation for more detailed information on configuring MWG. In the MWG UI under **Appliances > (this appliance) > Proxies (HTTP(S), FTP, SOCKS, ICAP...)**:

Web Proxy Service	User Input
NETWORK SETUP	Select **Proxy (Optional WCCP)**.
ADVANCED OUTGOING CONNECTION SETTINGS	**Check** the IP spoofing options. This will enable MWG to egress with the client's IP address.
OUTBOUND SOURCE IP LIST	Enter an IP address here in the same subnet as MWG's outbound interface (ex. 198.19.96.132). The SNAT policy will change data plane (client) traffic to this source address as it leaves the MWG.
HTTP PROXY	Click the plus sign to create a new HTTP proxy. Assign an appropriate listener IP address and port. This is the IP address and port that BIG-IP SSL Orchestrator will target. It is appropriate here to use a wildcard IP address (0.0.0.0) and the standard MWG explicit proxy port (**9090**). Click the OK button to save.

Click **Save Changes** in top right of the MWG UI to commit the changes.

To SNAT or not to SNAT

BIG-IP SSL Orchestrator is able to dynamically service chain traffic flows by a unique process of active "signaling". As a packet leaves the BIG-IP system for a security device, a marker is created, so that when the traffic returns, the marker restores context and continues moving the packets through the security stack. If a security device attempts its own external connection, that device-initiated traffic would not be signaled, thus the BIG-IP SSL Orchestrator device isolation posture would prevent it from traversing the inspection zone.

Should a security device require its own external connectivity then, it would be necessary to create a separate "control" channel on the BIG-IP system. This control channel is described more completely in the BIG-IP SSL Orchestrator configuration section below but is essentially a virtual server attached to the security device's from-service VLAN, listening on the device's outbound source IP address. Thus, to differentiate between decrypted client-server traffic that must continue through the service chain, and device-initiated traffic that must egress directly, minimally the source IP addresses must be different. MWG supports a feature called "IP spoofing," that when enabled retains the client's IP address on egress. It is therefore highly recommended to enable IP spoofing, so that only device-initiated traffic will source from the MWG's outbound IP address and be captured by the separate control channel virtual server, while normal client-server traffic sources from the client's IP address and flows through the service chain. It also benefits upstream security devices (devices later in the service chain) if they're able to see the client's real IP address.

Note: If IP spoofing is disabled, all traffic egresses MWG on the device's outbound interface IP address, which would then require a separate SNAT policy to translate client-server traffic to an alternate source IP to create the separate traffic patterns. As this is not recommended, instructions for creating this SNAT policy are not included here.

- In MWG **prior to 8.2**, IP spoofing functions natively by simply enabling the IP spoofing options in the "Proxies (HTTP(S), FTP, SOCKS, ICAP...)" section, under "Advanced Outgoing Connection Settings". You can then skip the following steps and move directly to BIG-IP SSL Orchestrator configuration.
- In MWG versions **8.2 and later**, IP spoofing must be enabled, but also requires a SNAT policy. The following steps detail that process.

Configure a SNAT policy

If running a version of MWG **prior to 8.2**, IP spoofing (egressing with the client's real IP address) requires both the IP spoofing settings enabled and a separate SNAT policy.

To create the SNAT policy, under the **Policy** section in the MWG UI, navigate to Common Rules and click the **Add Rule**... button.

SNAT Policy	User Input
NAME	Provide a unique name for this policy.
RULE CRITERIA	Select to apply this rule **Always**.
ACTION	Select **Continue**.

EVENTS	Click the Add button, and then Event... and create the following rule: `Enable Outbound Source IP Override (Client.IP)` • **Add -> Event....** • Select **Enable Outbound Source IP Override (IP)**. • Click the **Parameters**... button. • Click the **Parameter Property** button. • Select **Client.IP** and click the **OK** button. • Click the **OK** button, then again to complete the rule. Click **Finished** to complete the SNAT rule.

MWG should now be ready to receive decrypted HTTP traffic on its explicit proxy listener IP address and port. Data plane (client) traffic will egress MWG with the client's true self IP address, and any MWG-initiated traffic will egress on the appliance's from-service interface IP address. It is now time to configure BIG-IP SSL Orchestrator to integrate this MWG service.

CONFIGURE BIG-IP SSL ORCHESTRATOR

MWG is configured as an HTTP service in BIG-IP SSL Orchestrator. This configuration will focus on the traditional outbound (forward proxy) use case. Once logged into the BIG-IP system, navigate to the BIG-IP SSL Orchestrator menu, and review the environment.

Create the BIG-IP SSL Orchestrator deployment through Guided Configuration

The BIG-IP SSL Orchestrator Guided Configuration (GC) presents an entirely new and streamlined user experience. This workflow-based architecture provides intuitive, re-entrant configuration steps tailored to the selected topology.

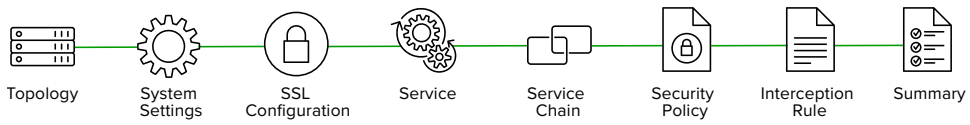

Figure 5: GC workflow

The following steps will walk through the GC to build a simple transparent forward proxy.

Initialization

If this is the first-time accessing BIG-IP SSL Orchestrator in a new BIG-IP build, upon first access, the GC will automatically load and deploy the built-in BIG-IP SSL Orchestrator package.

Configuration review and prerequisites

Take a moment to review the topology options and workflow configuration steps involved. Optionally satisfy any of the **DNS**, **NTP**, and **Route** prerequisites from this page. Keep in mind, however, that aside from NTP, the BIG-IP SSL Orchestrator GC will provide an opportunity to define DNS and route settings later in the workflow. No other configurations are required on this page, so click **Next**.

Figure 6: L3 Outbound GC

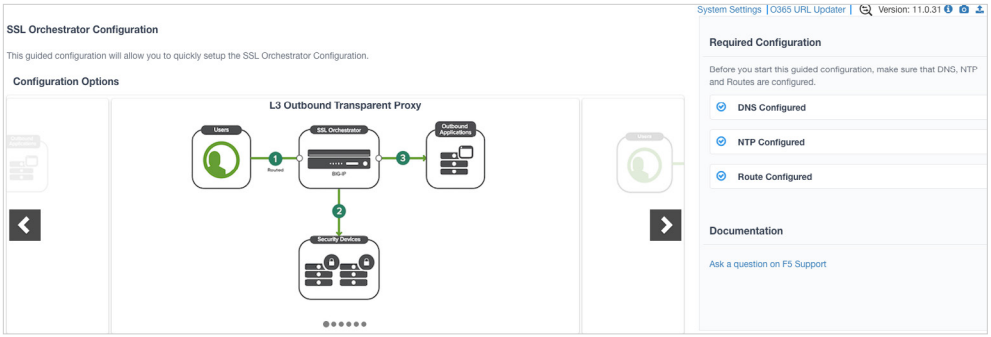

Topology properties

BIG-IP SSL Orchestrator creates discrete configurations based on the selected topology. An explicit forward proxy topology will ultimately create an explicit proxy listener and its relying transparent proxy tunnel. If a subsequent transparent forward proxy topology is configured, it will not overlap the existing explicit proxy objects. The Topology Properties page provides the following options:

Topology Properties	User Input
NAME	Enter a **Name** for the BIG-IP SSL Orchestrator deployment.
DESCRIPTION	Enter a **Description** for this BIG-IP SSL Orchestrator deployment.
CIPHER TYPE	The Protocol option presents four protocol types: • **TCP:** This option creates a single TCP wildcard interception rule for the L3 Inbound, L3 Outbound, and L3 Explicit Proxy topologies. SSL/TLS are primarily used with TCP protocols. A TCP topology will create the necessary architecture to decrypt and re-encrypt traffic flows. • **UDP:** This option creates a single UDP wildcard interception rule for L3 Inbound and L3 Outbound topologies. A UDP topology does not perform decryption but can service chain UDP traffic. • **Other:** This option creates a single any-protocol wildcard (all non-TCP/non-UDP traffic) interception rule for L3 Inbound and L3 Outbound topologies. A non-TCP/non-UDP topology does not perform decryption or service chaining. It creates a routed path for non-TCP/non-UDP traffic flows. • **Any:** This option creates the TCP, UDP, and non-TCP/UDP interception rules for outbound traffic flows. Select **TCP** to support decrypted traffic flows through the inline MWG.
IP FAMILY	Specify whether you want this configuration to support **IPv4** addresses or **IPv6** addresses.

BIG-IP SSL ORCHESTRATOR TOPOLOGIES	The BIG-IP SSL Orchestrator Topologies option page presents these six topologies:. • **L3 explicit proxy:** This is the traditional explicit forward proxy. • **L3 outbound:** This is the traditional transparent forward proxy. • **L3 inbound:** This is a reverse proxy configuration. • **L2 inbound:** The layer 2 topology options insert BIG-IP SSL Orchestrator as a bump-in-the-wire in an existing routed path, where BIG-IP SSL Orchestrator presents no IP addresses on its outer edges. The L2 Inbound topology provides a transparent path for inbound traffic flows. • **L2 outbound:** The layer 2 topology options insert BIG-IP SSL Orchestrator as a bump-in-the-wire in an existing routed path, where BIG-IP SSL Orchestrator presents no IP addresses on its outer edges. The L2 Outbound topology provides a transparent path for outbound traffic flows. • **Existing application:** This topology is designed to work with existing BIG-IP LTM applications. Whereas the L3 Inbound topology provides an inbound gateway function for BIG-IP SSL Orchestrator, Existing Application works with BIG-IP LTM virtual servers that already perform their own SSL/TLS handling and client-server traffic management. The Existing Application workflow proceeds directly to service creation and security policy definition, then exits with a BIG-IP SSL Orchestrator-type access policy and per-request policy that can easily be consumed by a BIG-IP LTM virtual server. Select L3 Outbound (transparent proxy) or L3 Explicit Proxy to support decrypted forward proxy traffic flows through the MWG.

Click **Save & Next**.

SSL configurations

This page defines the specific SSL/TLS settings for the selected topology—in this case a forward proxy—and controls both client-side and server-side SSL/TLS options. If existing SSL/TLS settings are available (from a previous workflow), it can be selected and re-used. Otherwise, the SSL Configurations page creates new SSL/TLS settings for this workflow.

SSL Configurations	User Input
SSL/TLS PROFILE	
NAME	Enter a **Name** for the SSL/TLS profile.
DESCRIPTION	Enter a **Description** for this SSL/TLS profile.
CLIENT-SIDE SSL/TLS	
CIPHER TYPE	Cipher type can be a Cipher Group or Cipher String. If the former, select a previously-defined cipher group (from Local Traffic – Ciphers – Groups). If the latter, enter a cipher string that appropriately represents the client-side SSL/TLS requirement. For most environments, **DEFAULT** is optimal

CERTIFICATE KEY CHAIN	The certificate key chain represents the certificate and private key used as the “template” for forged server certificates. While re-issuing server certificates on-the-fly is generally easy, private key creation tends to be a CPU-intensive operation. For that reason, the underlying SSL/TLS forward proxy engine forges server certificates from a single defined private key. This setting allows customers to apply their own template private key, and optionally store that key in a FIPS-certified HSM for additional protection. The built-in “default” certificate and private key uses 2K RSA and is generated from scratch when the BIG-IP system is installed. Click **Add**, select **default.crt** and **default.key**, and click **Done**.
CA CERTIFICATE KEY CHAIN	An SSL/TLS forward proxy must re-issue (“forge”) remote server certificate to local clients using a local CA certificate, and local clients must trust this local CA. This setting defines the local CA certificate and private key used to perform the forging operation. Assuming a local CA certificate and key were imported at the beginning of these steps, select them here. *Note: SSL/TLS settings minimally require RSA-based template and CA certificates but can also support Elliptic Curve (EC) certificates. In this case, BIG-IP SSL Orchestrator would re-issue (forge) an EC certificate to the client if the SSL/TLS handshake negotiated an ECDHE_ECDSA cipher. To enable EC forging support, add both an EC template certificate and key and EC CA certificate and key.*
[ADVANCED] BYPASS ON HANDSHAKE ALERT	This setting allows the underlying SSL/TLS forward proxy process to bypass SSL/TLS decryption if an SSL/TLS handshake error is detected on the server-side. It is recommended to leave this **disabled**.
[ADVANCED] BYPASS ON CLIENT CERTIFICATE FAILURE	This setting allows the underlying SSL/TLS forward proxy process to bypass SSL/TLS decryption if it detects a Certificate request message from the server, as in when a server requires mutual certificate authentication. It is recommended to leave this **disabled**.
Note: The above two Bypass options can create a security vulnerability. If a colluding client and server can force an SSL/TLS handshake error, or force client certificate authentication, they can effectively bypass SSL/TLS inspection. It is recommended that these settings be left disabled.	
SERVER-SIDE SSL/TLS	
CIPHER TYPE	Cipher type can be a Cipher Group or Cipher String. If the former, select a previously-defined cipher group (from Local Traffic – Ciphers – Groups). If the latter, enter a cipher string that appropriately represents the server-side SSL/TLS requirement. For most environments, **DEFAULT** is optimal.
TRUSTED CA	Browser vendors routinely update the CA certificate stores in their products to keep up with industry security trends, and to account for new and revoked CAs. In the SSL/TLS forward proxy use case, however, the SSL/TLS visibility product now performs all server-side certificate validation, in lieu of the client browser, and should therefore do its best to maintain the same industry security trends. BIG-IP ships with a CA certificate bundle that maintains a list of CA certificates common to the browser vendors. However, a more comprehensive bundle can be obtained from the F5 Downloads site. It is otherwise safe to select the built-in **ca-bundle.crt**.

[ADVANCED] EXPIRE CERTIFICATE RESPONSE	BIG-IP SSL Orchestrator performs validation on remote server certificates and can control what happens if it receives an expired server certificate. The options are drop, which simply drops the traffic, and **ignore**, which mirrors an expired forged certificate to the client. The default and recommended behavior for forward proxy is to **drop** traffic on an expired certificate.
[ADVANCED] UNTRUSTED CA	BIG-IP SSL Orchestrator performs validation on remote server certificates and can control what happens if it receives an untrusted server certificate, based on the Trusted CA bundle. The options are drop, which simply drops the traffic, and **ignore**, which allows the traffic and forges a good certificate to the client. The default and recommended behavior for forward proxy is to **drop** traffic on an untrusted certificate.
[ADVANCED] OCSP	This setting selects an existing or can create a new Online Certificate Status Protocol (OCSP) profile for server-side OCSP and OCSP stapling. With this enabled, if a client issues a Status_Request message in its ClientHello message (an indication that it supports OCSP stapling), BIG-IP SSL Orchestrator will issue a corresponding Status_Request message in its server-side SSL/TLS handshake. BIG-IP SSL Orchestrator will then forge the returned OCSP stapling response back to the client. If the server does not respond with a staple but contains an Authority Info Access (AIA) field that points to an OCSP responder URL, BIG-IP SSL Orchestrator will perform a separate OCSP request. The returned status is then mirrored in the stapled client-side SSL/TLS handshake.
[ADVANCED] CRL	This setting selects an existing or can create a new CRL profile for server-side Certificate Revocation List (CRL) validation. With this enabled, BIG-IP SSL Orchestrator attempts to match server certificates to locally cached CRLs.

Click **Save & Next**.

Services

The Services List page is used to define security services that attach to BIG-IP SSL Orchestrator. The GC includes a services catalog that contains common product integrations. Each icon represents a security product integration deployed as one of the five basic service types. The services catalog also provides “generic” security services if your security service is not included in the catalog. Depending on screen resolution, it may be necessary to scroll down to see additional services.

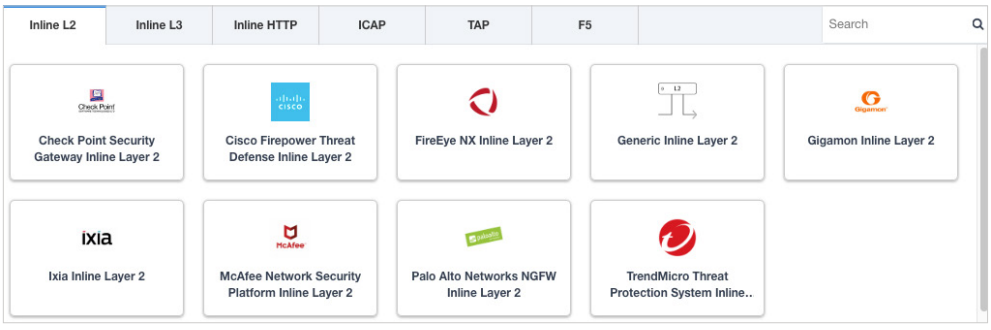

Figure 7: Services catalog

To define the MWG service, select it in the services catalog and click **Add**, or simply double-click the icon.

Services	User Input
NAME	Provide a unique name to this service (example “MWG”).
AUTO MANAGE ADDRESSES	When enabled the Auto Manage Addresses setting provides a set of unique, non-overlapping, non-routable IP addresses to be used by the security service. If disabled, the To and From IP addresses must be configured manually. It is recommended to leave this option **enabled (checked)**.
Note: In environments where BIG-IP SSL Orchestrator is introduced to existing security devices, it is a natural tendency to not want to have to move these devices. And while BIG-IP SSL Orchestrator certainly allows it, by not moving the security devices into BIG-IP SSL Orchestrator-protected enclaves, customers run the risk of exposing sensitive decrypted traffic, unintentionally, to other devices that may be connected to these existing networks. It is therefore highly recommended, and a security recommended practice, to remove BIG-IP SSL Orchestrator-integrated security devices from existing networks and place them entirely within the isolated enclave created and maintained by BIG-IP SSL Orchestrator.	
PROXY TYPE	This defines the proxy mode that the inline HTTP service is in. If in explicit proxy mode, the configuration will request the service’s listening IP address and proxy port. If in transparent proxy mode, the configuration will simply request the service’s inbound interface IP address. To support MWG explicit proxy mode authentication, configure this service as **Explicit**.
TO SERVICE CONFIGURATION	With the Auto Manage Addresses option enabled, this IP address will be pre-defined, therefore the to-service interface of the service must match this IP subnet. BIG-IP SSL Orchestrator uses an RFC2544 internal, non-routable address space (198.19.x.y/19). With the Auto Manage Addresses option disabled, the IP addresses must be defined manually. This assigned address specifies the BIG-IP VLAN self-IP from which traffic will be flowing to the service. • **To Service:** With the Auto Manage Addresses option enabled, this IP address will be pre-defined, therefore the inbound side of the service must match this IP subnet. With the Auto Manage Addresses option disabled, the IP address must be defined manually. • **VLAN:** Select the Create New option, provide a unique name (ex. **MWG_in**), select the BIG-IP interface connecting to the inbound side of the service, and add a VLAN tag value if required.
SERVICE DOWN ACTION	BIG-IP SSL Orchestrator also natively monitors the load balanced pool of security devices, and if all pool members fail, can actively bypass this service (**Ignore**), or stop all traffic (**Reset, Drop**).
SECURITY DEVICES	An inline HTTP service may be defined as the load balances set of multiple devices on the same IP subnets. Minimally one device IP address must be defined. • When the **Proxy Type** option is set to Explicit, this setting must specify the listening IP address and port of the HTTP proxy device. • When the **Proxy Type** option is set to Transparent, this setting must specify the to-service IP address of the HTTP proxy device. Click **Add**, enter the service’s inbound-side IP Address, the port value if explicit, then click **Done**.

FROM SERVICE CONFIGURATION	With the Auto Manage Addresses option enabled, this IP address will be pre-defined, therefore the from-service interface of the service must match this IP subnet. BIG-IP SSL Orchestrator uses an RFC2544 internal, non-routable address space (198.19.x.y/19). With the Auto Manage Addresses option disabled, the IP addresses must be defined manually. This assigned address specifies the F5 BIG-IP VLAN self-IP to which traffic will be flowing from the service back to the F5 BIG-IP system. This IP address should also be the gateway routing address on the layer 3 service to ensure all traffic is sent back to the F5 BIG-IP system. • **From Service:** With the Auto Manage Addresses option enabled, this IP address will be pre-defined, therefore the outbound side of the service must match this IP subnet. With the Auto Manage Addresses option disabled, the IP address must be defined manually. • **VLAN:** Select the Create New option, provide a unique name (ex. **MWG_out**), select the F5 BIG-IP interface connecting to the outbound side of the service, and add a VLAN tag value if required.
MANAGE SNAT SETTINGS	This setting allows BIG-IP SSL Orchestrator to SNAT traffic to an inline service. This is especially useful in a load-balanced BIG-IP SSL Orchestrator scaling configuration but is not required here. Leave it set to **None**.
AUTHENTICATION OFFLOAD	When an Access authentication profile is attached to an explicit forward proxy topology, this option will present the authenticated username value to the service as an X-Authenticated-User HTTP header. To **enable** delegate authentication to the MWG appliance, **enable** this option. MWG authentication is addressed later in this guide.
IRULES	BIG-IP SSL Orchestrator allows for the insertion of additional F5® iRules® logic at different points. An iRule defined at the service only affects traffic flowing across this service. It is important to understand, however, that these iRules must not be used to control traffic flow (ex. pools, nodes, virtuals, etc.) but rather should be used to view/modify application layer protocol traffic. For example, an iRule assigned here could be used to view and modify HTTP traffic flowing to/from the service. Additional iRules are not required, however, so leave this **empty**.

Click **Save**. When required services have been created, click **Save & Next**.

Service chains

Service chains are arbitrarily ordered lists of security devices. Based on environmental requirements, different service chains may contain different re-used sets of services, and different types of traffic can be assigned to different service chains. For example, HTTP traffic may need to go through all of the security services, while non-HTTP traffic goes through a subset, and traffic destined to a financial service URL can bypass decryption and still flow through a smaller set of security services.

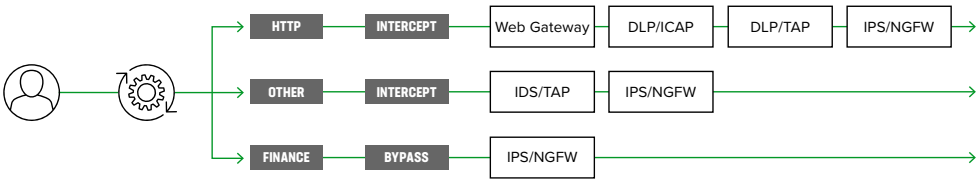

Figure 8: Different traffic flowing through chains of different security services

Click **Add** to create a new service chain containing all of the security services.

Service Chains	User Input
NAME	Provide a unique name to this service (ex. **"my_service_chain"**).
SERVICES	Select any number of desired services and move them into the **Selected Service Chain Order** column, optionally also ordering them as required. In this lab, select **all of the services**. Click **Save**.

Click **Save & Next**.

Security policy

Security policies are the set of rules that govern how traffic is processed in BIG-IP SSL Orchestrator. The "actions" a rule can take include:

- Whether or not to allow the traffic (Allow or Reject).
- Whether or not to decrypt the traffic (Intercept or Bypass).
- Which service chain (if any) to pass the traffic through.

The BIG-IP SSL Orchestrator GC presents an intuitive rule-based, drag-and-drop user interface for the definition of security policies. The security policy defines the set of traffic matching rules and corresponding actions to take on matches. The built-in security policy contains two rules:

- **Pinners_Rule**: Used to match URLs on a Pinners custom URL category, for SSL/TLS bypass. The Pinners custom URL category is a built-in category and contains a non-exhaustive list of sites known to use certificate pinning. You may need to add sites to this list based on specific business requirements.
- **All Traffic**: The default catch-all rule for any traffic that does not match other rules. By default, this rule allows and intercepts (decrypts) traffic but does not preselect any service chain. Optionally edit this rule to add a "default" service chain.

Figure 9: Configuring security policy

In the background, BIG-IP SSL Orchestrator maintains these security policies as visual per-request policies. If traffic processing is required that exceeds the capabilities of the rule-based user interface, the underlying per-request policy can be managed directly. By default, this rule allows and intercepts traffic but does not preselect any service chain.

Create any security rules as required, then click **Save & Next**.

Note: Once the per-request policy is manipulated, the rules-based user interface can no longer be used.

Interception rule

Interception rules are based on the selected topology and define the “listeners”, analogous to BIG-IP LTM virtual servers, which accept and process different types of traffic (ex. TCP, UDP, other). The resulting BIG-IP LTM virtual servers will bind the SSL/TLS settings, VLANs, IP addresses, and security policies created in the topology workflow. Again, note that security services are opaque within the protected inspection zone, thus the proxy mode deployed for the inline MWG is irrespective to the BIG-IP SSL Orchestrator topology mode. The following will demonstrate BIG-IP SSL Orchestrator as either a transparent or explicit forward proxy.

Interception Rule	User Input
COMMON SETTINGS	
SOURCE ADDRESS	The source address field provides a filter for incoming traffic based on source address and/or source subnet. It is usually appropriate to leave the default **0.0.0.0%0/0** setting applied to allow traffic from all addresses to be processed.
DESTINATION ADDRESS/MASK	The destination address/mask field provides a filter for incoming traffic based on destination address and/or destination subnet. As this is a transparent forward proxy configuration, it is appropriate to leave the default **0.0.0.0%0/0** setting applied to allow all outbound traffic to be processed.
INGRESS NETWORK—VLANS	This defines the VLANs through which traffic will enter. For a transparent forward proxy topology, this would be a client-side VLAN.
TRANSPARENT PROXY SETTINGS	
PORT	This defines a matching destination port for ingress traffic flows. It may be desirable to only process HTTP port 80 or HTTPS port 443 for example. The default 0 port pattern processes outbound traffic on any port.
SECURITY POLICY SETTINGS	This defines the security policy that will process traffic based on traffic matching rules. This will default to the previously created security policy.
EXPLICIT PROXY SETTINGS	
PROXY SERVER SETTINGS—IPV4 (OR IPV6) ADDRESS	This address is the IP address that clients will target to access external resources. This is typically done by setting the browser’s proxy server settings to this address, or by using Proxy Auto-Configuration (PAC) or WPAD scripts to point client user-agents at this IP address.
PROXY SERVER SETTINGS—PORT	The proxy service instance also requires a listening port. This is traditionally port 8080 or 3128.
PROXY SERVER SETTINGS—ACCESS PROFILE	In order to perform authentication on explicit forward proxy traffic, BIG-IP APM must be licensed and provisioned. Explicit forward proxy authentication is then defined within an “SWG-Explicit” access profile. Once an SWG-Explicit access profile is created, it can be selected here to enable authentication on this explicit forward proxy instance. Proxy authentication is covered in a later chapter.
Note: While visible for an explicit proxy topology, the port setting is grayed out and non-editable. The port for an explicit proxy is defined in the proxy server settings section of this page.	

Click **Save & Next**.

Egress settings

The Egress Setting page defines the topology-specific egress characteristics.

Egress Settings	User Input
MANAGE SNAT SETTINGS	Defines if and how SNAT is used for egress traffic.
GATEAWAYS	Defines the next hop route for traffic. For an outbound configuration, this is usually a next hop upstream router

Click **Save & Next**.

Summary

The summary page presents an expandable list of all of the workflow-configured objects. To expand the details for any given setting, click the corresponding arrow icon on the far right. To edit any given setting, click the corresponding pencil icon. Clicking the pencil icon will send the workflow back to the selected settings page.

When satisfied with the defined settings, click **Deploy**.

Upon successfully deploying the configuration, BIG-IP SSL Orchestrator will now display a **Configure** view.

Figure 10: Configure view

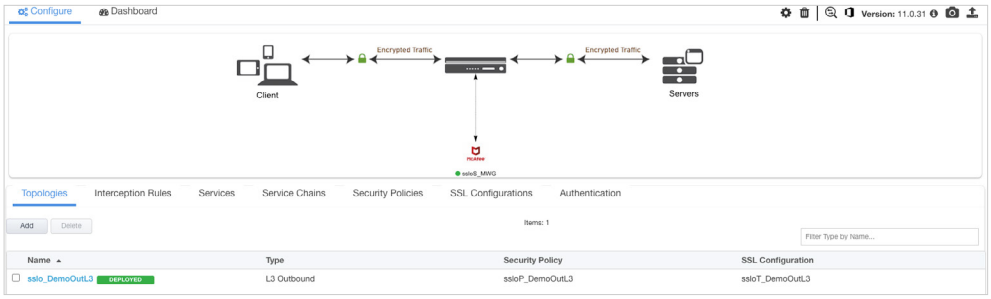

The **Interception Rules** tab may show one or two interception rules (listeners), depending on transparent or explicit proxy topology creation.

Figure 11: Interception Rules tab

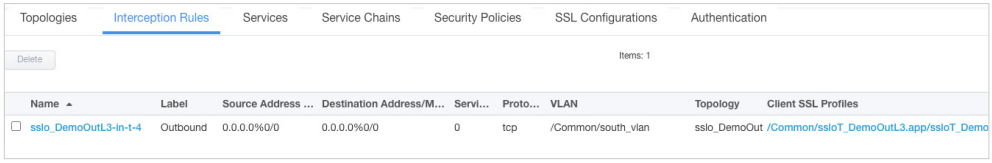

In the above,

- If a BIG-IP SSL Orchestrator L3 outbound (transparent proxy) topology was created:
  - The **-in-t-4** listener receives traffic from the client and performs all of the BIG-IP SSL Orchestrator functions (ex. decryption, service chaining).

- If a BIG-IP SSL Orchestrator explicit proxy topology was created:
  - The **-xp-4** listener is the explicit proxy service endpoint that receives traffic from the client. All HTTP and HTTPS traffic is tunneled through this listener.
  - The **-in-t-4** listener receives traffic from the -xp-4 tunnel and performs all of the BIG-IP SSL Orchestrator functions (ex. decryption, service chaining).

DNS query resolution (explicit proxy topology)

An explicit forward proxy performs DNS resolution on the client's behalf. For a BIG-IP SSL Orchestrator explicit proxy topology to work, you must also define DNS settings. Under the main BIG-IP SSL Orchestrator Configuration page, click on the gear icon in the top right to access System settings. This configuration step creates a DNS resolver object.

- Select Internet Authoritative Nameserver and enter Local/Private Forward Zones, or
- Select Local Forwarding Nameserver and enter Local DNS Nameservers.

Control channel virtual server

Previously, MWG was configured with IP spoofing, and optionally a SNAT policy to cause MWG to egress traffic on the client's IP address. This creates different traffic patterns for normal client-server traffic and device-initiated traffic so that a separate "control channel" virtual server can be created to catch device-initiated traffic and egress directly. In the BIG-IP UI, navigate to Local Traffic -> Virtual Servers and click the **Create** button.

SNAT Virtual Server	User Input
TYPE	Select **Performance (Layer 4)**.
SOURCE ADDRESS	Enter the MWG's from-service interface IP address here. This is the source address that the virtual server will filter on, to capture device-initiated traffic.
DESTINATION ADDRESS/MASK	Enter **0.0.0.0/0** to allow this virtual server to capture all MWG-initiated outbound traffic. You may optionally tighten this filter to specific IP addresses.
SERVICE PORT	Enter 0 to allow this virtual server to capture MWG-initiated traffic destined to any port. You may optionally tighten this filter to specific ports.
PROTOCOL	Select * **All Protocols**.
PROTOCOL PROFILE (CLIENT)	Select **fastL4**.
VLAN AND TUNNEL TRAFFIC	Enable only on the MWG from-service VLAN.
SOURCE ADDRESS TRANSLATION	Enable SNAT only if the BIG-IP SSL Orchestrator topology also defines egress SNAT.
ADDRESS TRANSLATION	This setting must be disabled (**unchecked**).

PORT TRANSLATION	This setting must be disabled (**unchecked**).
DEFAULT POOL	Create a new pool here and enter the same egress route IP address used in the BIG-IP SSL Orchestrator topology. If System Route was selected in the topology, do not assign a pool here.

Click **Finished** to complete this virtual server configuration. At this point, MWG will have access to Internet resources through this virtual server.

Note: This completes the configuration of BIG-IP SSL Orchestrator as a forward proxy. At this point an internal client should be able to browse out to external (Internet) resources, and decrypted traffic will flow across the security services.

Testing the Solution

You can test the deployed solution using the following options:

Server certificate test

To test an explicit forward proxy topology, configure a client's browser proxy settings to point to the listening IP address and port. Ensure that the client trusts the local issuing CA certificate. Open a browser from the client and attempt to access an external HTTPS resource. Once the page is loaded, observe the server certificate of that site, and take note of the certificate issuer, which should be the local issuing CA. If you have access to the client's command line shell and the cURL or wget utilities, you can simulate browser access using one of the following commands:

```
curl -vk –proxy [proxy IP:port] https://www.example.com
wget –no-check-certificate -e use_proxy=yes -e https_proxy=[proxy IP:port]
-dO – https://www.example.com
```

Both of these commands will display both the HTML server response, and the issuer of the server's certificate.

Decrypted traffic analysis on the BIG-IP system

Perform a tcpdump on the BIG-IP system to observe the decrypted clear text traffic. This confirms SSL/TLS interception by BIG-IP SSL Orchestrator.

```
Tcpdump –lnni [interface or VLAN name] -Xs0
```

As a function of adding a new service, the UI requires a name for each (source and destination) network. BIG-IP SSL Orchestrator will then create separate source and destination VLANs for inline security devices, and those VLANs will be encapsulated within separate application service paths. For example, given an inline HTTP service named "MWG" with its "From BIGIP VLAN" named "MWG_in", and its "To BIGIP VLAN" named "MWG_out", its corresponding BIG-IP VLANs would be accessible via the following syntax:

`ssloN_` + [network name] + `.app/ssloN_` + [network name]

Example:

`ssloN_MWG_in.app/ssloN_MWG_in`

`ssloN_MWG_out.app/ssloN_MWG_out`

A tcpdump on the source side VLAN of this MWG service would therefore look like this:

```
tcpdump -lnni ssloN_MWG_in.app/ssloN_MWG_in -Xs0
```

The security service VLANs and their corresponding application services are all visible from the BIG-IP UI under Network -> VLANs.

Decrypted traffic analysis on the MWG

From the MWG UI:

- Navigate to the Troubleshooting section and Packet Tracing.
- In the Command line parameters field, enter:

```
-s 0 -I any
```

This will capture traffic on all interfaces. To optionally capture on a specific interface, replace the word '**any**' above with the name of the interface.

- Click the **tcpdump** start button to start a capture, **tcpdump** stop the stop the capture, and then download the resulting capture to review (typically in Wireshark).
- It is also possible, if SSH access is enabled, to access the MWG console directly and issue **tcpdump** commands natively.

Proxy policy analysis on MWG

From the MWG UI:

- Navigate to the Troubleshooting section and Rule Tracing Central.
- Enter client IP address to be monitored and click Go button.
- Stop trace by clicking on the X button that "replaced" the Go button.
- More details can be found in the appropriate MWG Product Guide for your version.

Additional Considerations

There are a few other configuration concepts that you may need to explore when integrating a MWG with BIG-IP SSL Orchestrator.

AUTHENTICATION

The services attached to BIG-IP SSL Orchestrator are opaque to the external environment, thus they are protected, isolated, and do not interact outside of the internal connectivity with the BIG-IP system. If a MWG appliance requires authenticated user identity, BIG-IP SSL Orchestrator must be configured to perform the user authentication and pass identity information to MWG as a "delegate token". To do authentication, the BIG-IP system additionally requires BIG-IP APM activation. The following details this very simple configuration to enable delegate authentication to the MWG appliance.

Configure a BIG-IP APM access policy to authenticate explicit forward proxy users

Configuration of forward proxy client authentication is beyond the scope of this guide, except that explicit forward proxy authentication must use the SWG-Explicit profile type, and transparent forward proxy authentication uses the SWG-Transparent profile type for captive portal authentication.

- For explicit forward proxy authentication, create or edit an explicit proxy BIG-IP SSL Orchestrator topology and attach the SWG-Explicit access profile on the Interception Rules page of the topology workflow, under the Access Profile option.
- For transparent forward proxy (captive portal) authentication, refer to the BIG-IP SSL Orchestrator deployment guide for additional instructions.

Configure the MWG service definition to pass the delegate token

In the BIG-IP SSL Orchestrator UI, under Services, select to the edit the MWG service. Click to enable the Authentication Offload setting (checked). This option when enabled will send the authenticated user identity to the inline HTTP service as an X-Authenticated-User HTTP header.

Configure the MWG to consume the delegate token

Delegate token consumption is configured in a policy. In the MWG UI, in the Policy section, navigate to Common Rules, and click the **Add Rules**... button.

Delegate Token Policy	User Input
NAME	Provide a unique name for this policy.
RULE CRITERIA	Select **If the following criteria is matched** and click the **Add** button, and then **User/Group criteria**. Select the **Athentication.IsAuthenticated** option, **equals**, and **false**. Click the **OK** button.
ACTION	Select **Continue**.

EVENTS	Click the **Add** button, and then **Set Property Value**... and create the following rule:
	`Set Authentication.UserName = Header.Get("X-Authenticated-User")`
	`Set Authentication.IsAuthenticated = true`
	• **Add -> Set Property Value....**
	• Select **Authentication.UserName**.
	• Click the **Add**... button.
	• Click the **Parameter Property** button.
	• Select **Header.Get(String)** and click the **Parameters... button**.
	• In the Parameter Value field, enter **X-Authenticated-User** and click the **OK** button.
	• Click the OK button, then again to complete the rule.
	• **Add > Set Property Value....**
	• Select **Authentication.IsAuthenticated**.
	• In the Parameter value field, select **true**.
	• Click the **OK** button.

Click **Finished** to complete the delegate authentication rule.

Additional rules can optionally be created within the MWG to collect user group information from a username query to an external directory. Alternatively rules similar to the above can be used to process authenticated user groups in an X-Authenticated-Groups header populated by BIG-IP.

The MWG should now be ready to receive the user identity via the delegate X-Authenticated-User HTTP header. MWG may now use this user identity as if it had collected the value itself from direct challenge-response authentication.

IF CREATING SERVICE NETWORKS MANUALLY

Security recommended practice is to move security devices to BIG-IP SSL Orchestrator's protected network enclave. When doing so, BIG-IP SSL Orchestrator provides a set of internal network addresses that the security service can use. This is the **Auto Manage** setting in the service definition. In the event that you choose not to heed this security recommendation, or otherwise need to create alternate addressing for the inline security service, in the service definition, uncheck the Auto Manage option and perform the following steps.

Disable the Auto-Manage option and alter the **To Service Configuration** and **From Service Configuration** sections.

Egress Settings	User Input
TO SERVICE CONFIGURATION (SELECT CREATE NEW)	
NAME	Provide a unique local name (ex. MWG_in).
VLAN	Select **Create New** (or select an existing VLAN if one exists). • Select the correct inbound-side interface (BIG-IP to service). • Enter a VLAN tag value, as required.
NETWORK SETTINGS	• Enter the BIG-IP SSL Orchestrator inbound-side self-IP (ex. 198.19.96.7). • Enter the corresponding Netmask (ex. 255.255.255.128).
FROM SERVICE CONFIGURATION (SELECT CREATE NEW)	
NAME	Provide a unique local name (ex. MWG_out).
VLAN	Select **Create New** (or select an existing VLAN if one exists). • Select the correct inbound-side interface (service to BIG-IP). • Enter a VLAN tag value, as required.
NETWORK SETTINGS	• Enter the BIG-IP SSL Orchestrator outbound-side self-IP (ex. 198.19.96.245). • Enter the corresponding Netmask (ex. 255.255.255.128).

All other settings are the same.

DNS CACHING

In the above configurations, the MWG requires some amount of external connectivity, minimally to reach licensing and subscription update services, but also if an explicit proxy, to reach DNS. The control channel virtual server and MWG IP spoofing enables this direct connectivity. However, when BIG-IP SSL Orchestrator is also an explicit proxy, both devices require DNS, thus requiring two queries for each site accessed.

If BIG-IP SSL Orchestrator is licensed as an add-on to base BIG-IP LTM, the above DNS requirement can be further optimized using the F5® BIG-IP® DNS license. With this license, BIG-IP can act as a DNS cache for both BIG-IP SSL Orchestrator and for MWG. As traffic passes through the BIG-IP SSL Orchestrator explicit proxy, a DNS query is performed through the cache. If no record exists, external DNS is consulted, and a value returned and cached. When MWG performs its request, the record is immediately served from the same cache. To enable DNS caching, first ensure that the **BIG-IP DNS services license** is activated.

To create the DNS services cache configuration:

Create a DNS Cache configuration, in the **BIG-IP UI**, under **DNS > Caches > Cache List,** click **Create...**

DNS Cache	User Input
NAME	Provide a unique name.
RESOLVER TYPE	Select **Transparent (None)**.

All other settings here can be changed as needed.

Create a DNS profile, in the BIG-IP UI, under **Local Traffic** > **Profiles** > **Services** > **DNS**, click **Create...**

DNS Profile	User Input
NAME	Provide a unique name.
DNS CACHE	Select **Enabled**.
DNS CACHE NAME	Select the previously created DNS Cache.

Create a DNS Proxy virtual server. To allow access to the MWG, this virtual server should be on an address local to either the to-service or from-service VLANs defined for the MWG. In the BIG-IP UI, under **Local Traffic** > **Virtual Servers**, click **Create...**

DNS Cache Virtual Server	User Input
NAME	Provide a unique name.
TYPE	Select **Standard**.
SOURCE ADDRESS	Select **0.0.0.0/0**.
DESTINATION ADDRESS/MASK	Use an IP address here that is in the same subnet as the MWG's to-service or from-service VLAN.
SERVICE PORT	Enter port **53** here.
PROTOCOL	Select **UDP**.
PROTOCOL PROFILE (CLIENT)	Select **UDP**.
DNS PROFILE	Select the previously create DNS profile.
VLANS AND TUNNEL TRAFFIC	Select the corresponding MWG to-service or from-service VLAN.
SOURCE ADDRESS TRANSLATION	Enable SNAT as required to allow access to external DNS.
ADDRESS TRANSLATION	Check to enable.
PORT TRANSLATION	Check to enable.
DEFAULT POOL	Click the plus side (+) to create a new pool that points to an external DNS resource (ex. 8.8.8.8) or select an existing DNS services pool if one exists.

Finally, configure the MWG to point to this virtual server IP address for DNS queries. If BIG-IP SSL Orchestrator is configured as an explicit proxy, configure its DNS Local Forwarding Nameserver settings to point to this same virtual server IP address.

TRANSPARENT PROXY PASS-THROUGH AUTHENTICATION

MWG provides transparent proxy authentication via an HTTP redirect mechanism to an on-box authentication service, which supports multiple forms of authentication. The mechanism normally works by redirecting the user to an authentication service URL on the MWG appliance, where the user authenticates and is then redirected back through the transparent proxy.

A summary of the transparent proxy authentication traffic flow:

Flow	Example
THE CLIENT ISSUES A NORMAL HTTP REQUEST THROUGH THE TRANSPARENT PROXY.	GET /foo Host: www.example.com
IN THE ABSENCE OF AN AUTHENTICATED SESSION, THE PROXY REDIRECTS THE USER TO A SEPARATE AUTHENTICATION SERVICE URL.	HTTP/1.1 302 Proxy Redirect Location: https://mwg.auth.url:port/mwg-internal/<uid>/plugin?target=Auth&reason=Auth&ClientID=<uid>&url=<value>&rnd=<uid>
THE CLIENT IS CHALLENGED AND PRESENTS CREDENTIALS TO THE AUTHENTICATION SERVICE.	401-based authentication Basic, NTLM or Kerberos (or something else)
THE CLIENT REDIRECTS BACK TO THE ORIGINAL URL.	GET /foo Host: www.example.com

This logical flow presents two challenges when MWG is inside the BIG-IP SSL Orchestrator inspection zone:

- All traffic to the MWG is unencrypted HTTP, therefore the Original-URL in the authentication service redirect always includes an HTTP:// URL. To address this challenge, a simple iRule is required to rewrite the Original-URL value in the authentication service redirect as it leaves the MWG.
- The authentication service URL is defined within the MWG configuration and requires DNS to resolve to an IP. However, the IP address of the MWG inside the BIG-IP SSL Orchestrator inspection is non-routable and inaccessible to the client. To address this challenge, a separate F5 virtual server and pool are required to direct traffic to the MWG authentication service. Authentication service traffic should DNS resolve to this client-accessible IP on the F5 system, which then passes the traffic to the internal service.

The iRule and other configuration steps are detailed below.

MWG: Configure a transparent proxy with authentication

This guide does not expand on all the different ways to configure transparent proxy authentication on MWG. Please see the official McAfee documentation for detailed options.

The following steps approximate the configuration of a transparent proxy:

- In the Configuration utility, under **Appliances > <appliance> Proxies HTTP(S), FTP, SOCKS, ICAP...**, select to enable the "**L2 Transparent**" proxy option. Set port redirects for port 80 and 443 to the configured explicit proxy port (9090 by default). In the Advanced Outgoing Connection Settings section, also enable **IP spoofing (HTTP, HTTPS, FTP)**.
- In the Policy utility, add a new Top Level Rule Set. Select to **Import rule set from Rule Set Library**. In the Rule Set Library dialog, select **Authentication > Authentication Server**. If there are conflicts, select Auto-Solve Conflicts then Solve by Referring to Existing Objects and click OK.
    - Select the **Authentication Server** top-level rule set and click on Unlock the view. When prompted are you sure, select Yes.
    - Select the **Check for Valid Authentication Session** ruleset, then edit the **Redirect Clients That do Not Have a Valid Session to the Authentication Server** rule.
      In the Rule Criteria section, edit the first rule, then edit the Settings of that **IP Authentication Server** rule:
        - Authentication method: Authentication server.
        - Authentication server URL: Change this value to reflect the actual URL the user will be redirected to. This will be an external URL hosted on a BIG-IP virtual server (ex. https://login.f5labs.com).
    - Select the **Authentication Server** ruleset within the Authentication Server Top Level Ruleset.
        - Select the **Authenticate User Against User Database** Rule and right click to edit.
        - Edit the Rule Name as desired (example: Authenticate User Against Active Directory).
        - Click Next.
        - Edit the Rule Criteria by double clicking on the only Criteria listed.
        - The **Authentication.Authenticate** property will be pre-selected as the property. Use the pull-down Settings: button within the property to **Add Authentication Settings** or select an existing set of authentication settings.

- Name the new set of authentication settings (example NTLM for Authentication Server).
- Choose your **Authentication Method**.
- Set remaining settings as desired for your chosen method.
- Click OK.
- Make sure the selected operator is still equals and the Compare with is still False.
- Click OK.
- Check that the Action is still **Authenticate** with Settings Default.
- Click Finish.

◦ Drag the Top-Level Ruleset to the appropriate place in your rule tree.

• Submit and commit changes.

BIG-IP SSL Orchestrator: Create a pool to the MWG authentication service

With HTTPS authentication disabled in the MWG authentication realm settings, traffic to the authentication service defaults to port 80.

Note: The authentication service pool must point to the same IP addresses defined in the BIG-IP SSL Orchestrator service, but on port 80.

BIG-IP SSL Orchestrator: Create a client SSL/TLS profile for the MWG authentication service

Authentication traffic should still flow across an encrypted channel, so client SSL/TLS must be configured on the F5 system to decrypt to the MWG. The certificate and private key defined in the F5 client SSL/TLS profile should match the URL defined in the MWG redirect hostname.

BIG-IP SSL Orchestrator: Create a source address persistence profile

To enable authentication service requests and regular decrypted traffic to flow to the same MWG service for a single user, a persistence profile is required that maps the source address of the client. Create a Source Address Affinity persistence profile, with "**Match Across Virtual Servers**" enabled.

BIG-IP SSL Orchestrator: Create a virtual server to the MWG authentication service

The F5 virtual server must minimally include:

- **Destination Address/Mask**—an IP address that matches the DNS resolution of the Redirect Hostname value.
- **Port**—port 443.
- **Client SSL/TLS profile**—the previously created client SSL/TLS profile.
- **VLANs and Tunnels**—the client-side VLAN.
- **Address Translation**—enabled.

- **Port Translation**—enabled.
- **Default Pool**—the previously-created pool.
- **Default Persistent Profile**—previously-created source address affinity profile.

BIG-IP SSL Orchestrator: Create an iRule for the MWG authentication redirects

The iRule must be attached to the MWG service to catch and rewrite the authentication redirect responses coming from the MWG service. In the BIG-IP SSL Orchestrator UI, under Services, click to edit the MWG service. Under the Resources section at the bottom of the MWG service properties page, select the created iRule, then click Save & Next. Deploy to save this configuration.

```
when HTTP_RESPONSE {
    ## Trigger URL rewrite if an mwg redirect that contains a url
query string parameter
    if { ( [HTTP::is_redirect] ) and ( [HTTP::header Location] contains
"/mwg-internal/" ) and ( [HTTP::header Location] contains "&url=" ) } {
        ## Evaluate original external request to SSLO (http or https?)
        sharedvar ctx
        if { ( [info exists ctx(ptcl)] ) and ( $ctx(ptcl) eq "https" )
} {
            ## Replace http:// in auth URL with https://
            set url_orig [findstr [HTTP::header Location] "&url=" 5
"&"]
            set url_dec [b64decode [URI::decode $url_orig]]
            set url_rew [string map {"http://" "https://"} $url_dec]
            set url_enc [URI::encode [b64encode $url_rew]]
            set loc_new [string map [list ${url_orig} ${url_enc}]
[HTTP::header Location]]

            HTTP::header replace Location ${loc_new}
        }
    }
}
```

BIG-IP SSL Orchestrator: Attach the new iRule to the MWG service

Navigate to the BIG-IP SSL Orchestrator services menu and click on the MWG service. Click on the pencil icon to the right of the Service section to edit the service. At the bottom of the MWG service configuration, move the above authentication rule to the Selected block. Click **Save & Next**. Click **Save & Next** again at the Service Chain list, then click **Deploy**.