Configuring F5 for Air Gap Egress Inspection with SSL Intercept

Welcome to the F5® deployment guide for configuring the BIG-IP® system for air gap SSL inspection. This document contains guidance on configuring the BIG-IP system to act as a forward proxy, decrypting outbound HTTPS traffic so it can be inspected by a security device, which then forwards the traffic to BIG-IP for re-encryption and delivery to the destination.

This guide provides instructions on configuring the BIG-IP system version 11.4 and later using an iApp™ application template to simplify deployment and maintenance (the SSL category bypass requires 11.5 or later and a URL Filtering subscription). There is also an appendix with manual configuration tables for users who prefer to create each individual object.

Why F5?

SSL Visibility

SSL termination is resource-intensive. F5 BIG-IP devices include dedicated hardware processors specializing in SSL processing. In both inbound and outbound deployment scenarios, using F5 SSL Intercept solution provides uncompromising visibility into SSL traffic.

The proliferation of websites now leveraging SSL encryption to protect users poses a challenge to security sensor pools in their mission to eliminate malware and attacks for outbound application requests. With the BIG-IP LTM, SSL Intercept can be leveraged to provide full visibility into user traffic.

For those with policy and privacy concerns, SSL category bypass can be configured to not decrypt requests to sites with sensitive data.

Products and applicable versions

Product	Version
BIG-IP LTM, AFM	11.4 - 11.6 (11.5 and later if using SSL Forward Proxy bypass)
iApp Template Version	f5.airgap_egress.v1.0.0rc4
Deployment guide version	1.3 (see *Document Revision History on page 33*)

Important: *Make sure you are using the most recent version of this deployment guide, available at http://www.f5.com/pdf/deployment-guides/f5-airgap-dg.pdf.*

To provide feedback on this deployment guide or other F5 solution documents, contact us at *solutionsfeedback@f5.com*

Contents

What is F5 iApp?

New to BIG-IP version 11, F5 iApp is a powerful new set of features in the BIG-IP system that provides a new way to architect application delivery in the data center, and it includes a holistic, application-centric view of how applications are managed and delivered inside, outside, and beyond the data center. The iApp template acts as the single-point interface for managing this configuration.

For more information on iApp, see the White Paper *F5 iApp: Moving Application Delivery Beyond the Network:* *http://www.f5.com/pdf/white-papers/f5-iapp-wp.pdf*.

Prerequisites and configuration notes

The following are general prerequisites and configuration notes for this guide:

- For this guide, the BIG-IP system ***must*** be running version 11.4 or later. This guide does not apply to previous versions.
- If you are running version 11.5 or later, have an active URL Filtering Subscription and have Secure Web Gateway (SWG) provisioned, you can select specific URL categories to bypass SSL filtering. See *Optional URL filtering on page 4* for more details. Contact your F5 sales representative for information on the URL Filtering Subscription.
- This guide describes two configuration scenarios: using a single BIG-IP device, and using two BIG-IP devices (an internal device and an external device). If you are using the iApp template in the two BIG-IP device scenario, you would run the iApp once on each BIG-IP device, once for receiving traffic from clients, and once for receiving traffic from a security device. See the Configuration example on this page for more information.
- If you are deploying a single BIG-IP LTM, you must have at least two VLANs configured on the system (one for receiving traffic from clients and one for receiving traffic from a security device). See the BIG-IP documentation for information on how to configure VLANs. See the following section for details on one and two device solutions.
- If you are deploying this configuration to forward traffic to an OSI layer 3-based security device, the security device(s) must have a default gateway defined to be the internal self IP address of the egress BIG-IP in a two-box deployment, or the egress network in a single-box deployment."
- For this configuration, you must have imported a certificate and key from a Certificate Authority which are trusted by your internal clients onto the BIG-IP system. To import certificates and keys, see **System > File Management > SSL Certificate List**. For specific instructions on importing certificates and keys, see the Help tab or the BIG-IP system documentation on *support.f5.com*. The importing process is not a part of the iApp template.
- This iApp configures a single wildcard ingress TCP virtual server for SSL detection. All encrypted, non-HTTP traffic (FTPS, SMTPS, etc.) traversing this virtual server will fail. To allow BIG-IP to pass traffic for these encrypted protocols, you must configure a separate virtual server for each service; to pass SMTPS traffic, for example, use the SMTP iApp template and deployment guide found here: *http://www.f5.com/pdf/deployment-guides/f5-smtp-dg.pdf*.

Configuration example

In this guide, we describe how to configure the BIG-IP system as an SSL Forward Proxy with Intercept Air Gap. This means that the BIG-IP system decrypts SSL traffic from internal clients, forwards the unencrypted traffic to a security device for inspection, and then re-encrypts the traffic in a way the client expects. There are two configuration options, depending on whether your implementation is using a single BIG-IP system, or an internal and external BIG-IP system.

The internal/external BIG-IP system scenario looks like the following diagram:

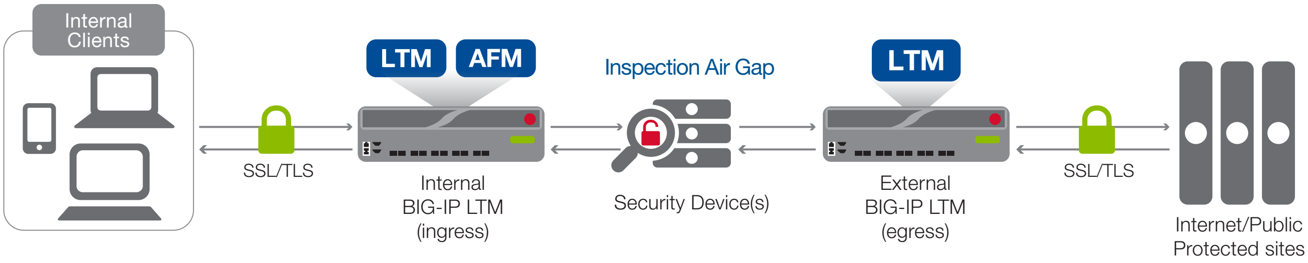

Figure 1: *Logical configuration diagram of the air gap configuration using an Internal and External BIG-IP system*

The traffic flow for this scenario is:

1. An internal client requests an encrypted site, and because of default route settings, the browser sends the request to the internal BIG-IP LTM.
2. The client initiates an SSL session with the internal BIG-IP LTM.
3. The BIG-IP LTM initiates a separate SSL session with the remote host the client requested (through the external BIG-IP LTM). The remote host sends its server certificate to the internal BIG-IP LTM as part of the negotiation.
4. The internal LTM generates a server certificate on-the-fly to match the properties of the remote host's server certificate and presents that to the client to complete the client-side SSL negotiation.
5. After completing the SSL handshake, the client sends its HTTP request. The internal BIG-IP LTM detects this request, disables server-side SSL, injects a special HTTP header, and sends the traffic to the next destination. This destination depends on how you are configuring the solution:
   - If you are configuring a transparent solution, the internal BIG-IP LTM sends the traffic to the self IP address of the External BIG-IP LTM. In this case, there is a security device between the two systems which can transparently inspect the traffic (see Figure 1).
   - If you are deploying this configuration to forward traffic to an OSI layer 3-based security device, the security device(s) must have a default gateway defined to be the internal self IP address of the egress BIG-IP in a two-box deployment, or the egress network in a single-box deployment.
6. The external LTM receives the request on its port 80 wildcard virtual server, detects the special HTTP header, applies a server SSL profile, and then sends the data to the destination address on port 443.

If you are deploying a single box scenario, the flow is largely the same, but the BIG-IP LTM must listen on separate VLANs for connections from the client devices and the security devices. This is a routed solution, and the security devices must have a default gateway that points back to the BIG-IP system on the proper VLAN.

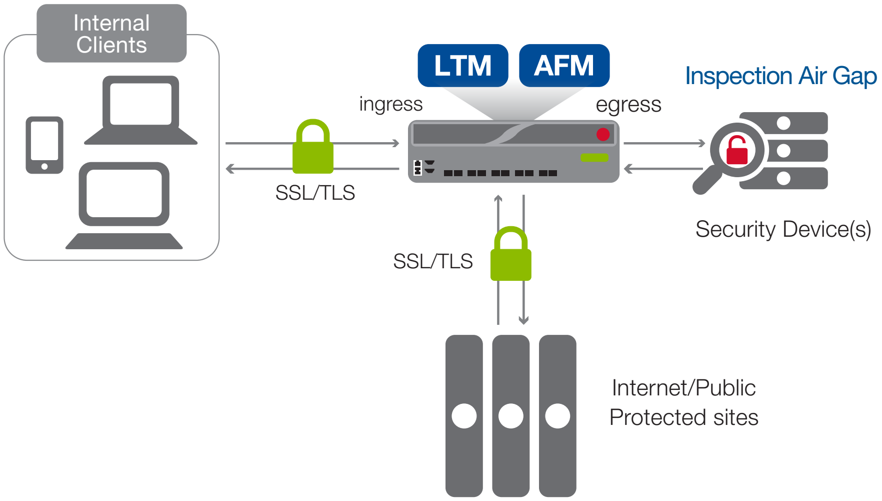

Figure 2: *Single box logical configuration example*

Optional URL filtering

If you have licenced URL filtering on your BIG-IP system (v11.5 and later), and have Secure Web Gateway (SWG) provisioned (it does not have to be licensed), you can add filtering to the implementation. This allows you to select specific URL categories that should bypass SSL decryption. Normally this is done for concerns over user privacy, or for categories that contain items that may rely on specific SSL certificates to be presented as part of a verification process (such as software update tools).

Configuring the BIG-IP system using the iApp template

Use this section if you plan on configuring the BIG-IP system using the Air Gap iApp template. If you plan to configure the BIG-IP system manually, see *Appendix A: Manual Configuration tables on page 21.*

Downloading and importing the Air Gap iApp template

The next task is to download and import the iApp template.

To download and import the iApp

1. Using a web browser, go to: *https://devcentral.f5.com/codeshare?sid=319* and download and extract the **f5.airgap_egress.v1.0.0rc4** (or newer) iApp template.
2. Log on to the BIG-IP system web-based Configuration utility.
3. On the Main tab, expand **iApp**, and then click **Templates**.
4. Click the **Import** button on the right side of the screen.
5. Click a check in the **Overwrite Existing Templates** box.
6. Click the **Browse** button, and then browse to the location you saved the iApp file.
7. Click the **Upload** button. The iApp is now available for use.

Starting the iApp template

To begin the iApp Template, use the following procedure.

To start the iApp template

1. Log on to the BIG-IP system.
2. On the Main tab, expand **iApp**, and then click **Application Services**.
3. Click **Create**. The Template Selection page opens.
4. In the **Name** box, type a name. In our example, we use **airgap-.**
5. From the **Template** list, select **f5.airgap_egress.v1.0.0rc4** (or newer if applicable).

Template Options

At the bottom of the Welcome section of the iApp template, you will find the following general questions.

1. ***Do you want to see inline help***
   Select whether you want to see informational and help messages inline throughout the template. If you are unsure, we recommend leaving the default, **Yes, show inline help**.
   Important and critical notes are always shown, no matter which selection you make.

   - **Yes, show inline help**
     Select this option to show inline help for most questions in the template.
   - **No, do not show inline help**
     Select this option if you do not want to see inline help. If you are familiar with this iApp template, or with the BIG-IP system in general, you can select this option to hide the inline help text.

2. ***Which configuration mode do you want to use?***
   Select whether you want to use F5 recommended settings, or have more granular, advanced options presented.

   - **Basic - Use F5's recommended settings**
     In basic configuration mode, options like load balancing method, parent profiles, and settings are all set automatically. The F5 recommended settings come as a result of extensive testing, so if you are unsure, choose Basic.

- **Advanced - Configure advanced options**
  In advanced configuration mode, you have more control over individual settings and objects, such the ability to restrict traffic to specific VLANs or attach iRules you have previously created to the application service. This option provides more flexibility for advanced users.

  Advanced options in the template are marked with the Advanced icon: **Advanced** . If you are using Basic/F5 recommended settings, you can skip the questions with this icon.

Virtual Server configuration

Use this section for information on the questions related to the virtual server, such as certificate and key information, and profiles.

1. *Where does this BIG-IP system reside in your network?*
   Select where you have deployed this BIG-IP system in your network topology. The BIG-IP system needs to know from where it will be receiving traffic, your selection determines the rest of the questions in the iApp.

   ➡ ***Note*** *The first two selections ("This LTM will receive ingress traffic from internal clients", and "This BIG-IP LTM will receive egress traffic from a security device") are meant to be used in a two BIG-IP device configuration. Using this scenario, you would run the iApp on each BIG-IP device, selecting the appropriate answer from this list for each device.*

   - If the BIG-IP LTM is receiving traffic from internal clients (the traffic ingress point coming from the client), continue with the first bullet on this page.

   - If the BIG-IP LTM is receiving traffic from a security device (the traffic egress point coming from the security device), jump ahead to *This BIG-IP LTM will receive egress traffic from a security device on page 11)*,

   - If the BIG-IP LTM is receiving traffic from both clients and a security device, jump head to *This BIG-IP LTM will receive ingress and egress traffic on different networks on page 13)*.

- **This BIG-IP LTM will receive ingress traffic from internal clients**
  Select this option if the BIG-IP system you are configuring is facing the internal network, and is going to be receiving traffic from internal clients.

  a. *Which type of forward proxy are you deploying at this time?*
  Choose whether you want to deploy the system as an Explicit Forward Proxy or a Transparent Forward Proxy. Selecting Explicit Proxy configures an HTTP proxy in explicit mode. In this scenario, browser clients must be specifically configured to use the proxy via system settings. Selecting Transparent Proxy configures a transparent forwarding proxy. In this scenario, browser clients passing through the BIG-IP system do not need to be configured with proxy settings. Note that this mode requires the traffic from the clients is routed to one of the BIG-IP Self IP addresses as part of its route out of the network.

    - **Transparent Proxy**
      Select this option if you want to configure the BIG-IP system as a transparent forward proxy. Continue with *b. Which Client SSL profile (with SSL Forward proxy enabled) do you want to use?* on the next page.

    - **Explicit Proxy**
      Select this option if you want to configure the system as an Explicit proxy. You must answer the following questions.

      a. *What IP address and port do you want to use for the proxy virtual server?*
      Type the IP address and port (both are required) for your explicit proxy instance. Only change the port if you have modified it from the default port, **3128**.

      b. *What is the FQDN of this proxy?*
      Type the fully qualified domain name of your explicit proxy.

      c. *Do you want the system to forward all name requests?*
      Chose whether you want the system to forward all name requests to a group of DNS servers, or if you want the BIG-IP system to resolve and cache the names directly (e.g. follow root hints). The default is to forward DNS name requests to a group of external resolvers.

        - **No, resolve all names directly**
          Choose this option if you want the BIG-IP system to resolve and cache the names directly. Continue with the next question.

        - **Yes, forward all requests**
          Choose this option if you want the system to forward all DNS requests. You must specify the servers to which you want to forward requests in the next question.

a. *Which DNS servers do you want to use for forwarding?*
Specify the IP address(es) of the DNS servers you will use to resolve external host names by this proxy instance. Click the Add button to include more servers. If you are using a different port than the default (53), make sure to modify the Port value.

b. *Which Client SSL profile (with SSL Forward proxy enabled) do you want to use?*
Specify whether you want the template to create a new Client SSL profile, or if you created a custom Client SSL profile that has SSL Forward Proxy enabled and uses a CA certificate trusted by your clients, you can select it from the list.

ⓘ *Important* *The CA certificate (and private key) used here is for issuing new server certificates. The CA certificate must have the Digital Signature and Certificate Signing key usage properties (at a minimum). We recommend using a subordinate CA certificate if available.*

- *Select the Client SSL profile you created from the list*
If you created a custom Client SSL profile for this implementation, select it from the list. This profile must contain a certificate and key from a Certificate Authority, and they must be trusted by your internal clients. For information on creating a custom Client SSL profile and importing certificates and keys, see the Help tab or the BIG-IP documentation on *support.f5.com*.

- **Create a new Client SSL profile**
Select this option if you want the iApp template to create a new SSL profile as a part of this application service. Although the iApp template creates the profile, it cannot import the certificate and key. You must have already imported a certificate and key from a Certificate Authority which are trusted by your internal clients onto this BIG-IP system. To import certificates and keys, see **System > File Management > SSL Certificate List**. For specific instructions, see the Help tab or the BIG-IP documentation on *support.f5.com*.

a. *Which trusted CA certificate do you want to use to issue server certificates for client-side connections?*
Select the CA certificate you imported onto the BIG-IP system for client authentication that is trusted by your internal clients.

⚠ *Warning* *Remember, this certificate must be issued by Certificate Authority, trusted by your internal clients., and have the Digital Signature and Certificate Signing key usage properties. We recommend using a subordinate CA certificate if available.*

b. *Which trusted CA private key do you want to use to issue server certificates for client-side connections?*
Select the SSL private key associated with the certificate you selected.

c. *Which hostnames would you like to bypass SSL interception?* **Advanced**
If you created a BIG-IP *Data Group* object containing host names that you want to bypass SSL interception, you can select it from the list. Only previously created Data Group objects appear in the list. Creating a Data Group is not a part of this template; if you want to create a Data Group, go to **Local Traffic > iRules > Data Group List**. For specific information, see the Help tab or the BIG-IP documentation.

- **Do not bypass hostnames**
Select this option if you do not want any hostnames to bypass SSL interception, or if you have not yet created a Data Group with hostnames.

- *Select the Data Group list you created with the hostnames you want to bypass*
Select the Data Group list you created for the hostnames you want to bypass.

d. *Which source IP addresses would you like to bypass SSL interception?* **Advanced**
If you created a BIG-IP *Data Group* object containing source IP address that you want to bypass SSL interception, you can select it from the list. Only previously created Data Group objects appear in the list. Creating a Data Group is not a part of this template; if you want to create a Data Group, go to **Local Traffic > iRules > Data Group List**. For specific information, see the Help tab or the BIG-IP documentation.

- **Do not bypass source IP**
Select this option if you do not want any source IP addresses to bypass SSL interception, or if you have not yet created a Data Group with source IP addresses.

- *Select the Data Group list you created with the source IP addresses you want to bypass*
Select the Data Group list you created for the source IP addresses you want to bypass.

e. *Which destination IP addresses would you like to bypass SSL interception?* **Advanced**
If you created a BIG-IP *Data Group* object containing destination IP addresses you want to bypass SSL interception, you can select it from the list. Only previously created Data Group objects appear in the list.

Creating a Data Group is not a part of this template; if you want to create a Data Group, go to **Local Traffic > iRules > Data Group List**. For specific information, see the Help tab or the BIG-IP documentation.

- **Do not bypass destination IP addresses**
  Select this option if you do not want any destination IP addresses to bypass SSL interception, or if you have not yet created a Data Group with destination IP addresses.
- *Select the Data Group list you created with the destination IP addresses you want to bypass*
  Select the Data Group list you created for the destination IP addresses you want to bypass.

f. *Which certificate bundle contains your Trusted Root CAs?* Advanced
Select the certificate bundle that contains your Trusted Root Certificate Authorities. For this question, you can leave the default if applicable, as it contains many of the most common Certificate Authorities. You can view the list at **System > File Management > SSL Certificate List > ca-bundle**, in the **Certificate Subject(s)** field.

g. *What action should be taken for an expired certificate?* Advanced
Choose the action you want the BIG-IP system to perform if the certificate has expired.

- **Drop**
  Select this option if the BIG-IP system should drop the connection from a client with an expired certificate.
- **Ignore**
  Select this option if the BIG-IP system should ignore the expired certificate and allow the connection.

h. *What action should be taken for an untrusted certificate?* Advanced
Choose the action you want the BIG-IP system to perform if the certificate is not trusted.

- **Drop**
  Select this option if the system should drop the connection from a client with an untrusted certificate.
- **Ignore**
  Select this option if the BIG-IP system should ignore the untrusted certificate and allow the connection.

i. *To which device(s) should this BIG-IP LTM forward decrypted outbound client traffic?*
Specify the IP address of each device to which the BIG-IP system should forward outbound client traffic. For the ingress device, the destination should be the self IP address of the egress device if the security device between ingress and egress points is OSI layer 2-based. If the device is OSI layer 3-based, this could be the IP address of the security device itself, or another device configured to route traffic to the security device.

Click **Add** to include more devices. The system creates load balancing pools with the addresses you specify here.

j. *Which HTTP profile do you want to use for client-side traffic?*
The HTTP profile contains settings for instructing the BIG-IP system how to handle HTTP traffic. Choose whether you want the iApp to create a new HTTP profile or if you have previously created an HTTP profile for this deployment.

Unless you have requirements for configuring specific HTTP settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic** >> **Profiles** : **Services : HTTP** to create a HTTP profile. To select any new profiles you create, you need to restart or reconfigure this template.

- *Select an existing HTTP profile from the list*
  If you already created an HTTP profile for this implementation, select it from the list.
- **Create a new HTTP profile (recommended)**
  Select this option to have the iApp to create a new HTTP profile.

k. *Do you want to restrict client traffic to specific VLANs?*
The BIG-IP system allows you to restrict client traffic to specific VLANs that are present on the system. This can provide an additional layer of security, as you can allow or deny traffic from the VLANs you choose. By default, all VLANs configured on the system are enabled. If you select to enable or disable traffic on specific VLANs, you must specify the VLANs in the next question. The VLAN objects must already be configured on this BIG-IP system before you can select them.

- **Enable traffic on all VLANs and Tunnels**
  Choose this option to allow traffic from all VLANs and Tunnels. If you select this option, the question asking about VLANs disappears. Continue with the next question.
- **Yes, enable traffic only on the VLANs I specify**
  Choose this option to restrict client traffic to specific VLANs that you choose in the following question. The system will accept client traffic from these VLANs, and deny traffic from all other VLANs on the system.

a. *<u>On which VLANs should traffic be enabled or disabled?</u>*
Use this section to specify the VLANs that accept client traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so click the VLANs and then use the Move buttons to adjust list membership.

➡ ***<u>Note:</u>*** *If you choose to allow traffic from certain VLANs, when additional VLANs are added to the BIG-IP system at a later time, this iApp configuration will deny traffic from these VLANs by default. To accept traffic from these VLANs, you must re-enter the template and add the VLAN(s).*

- **Yes, disable traffic only on the VLANs I specify**
Choose this option to deny client traffic from the specific VLANs that you choose in the following question. The system refuses client traffic from these VLANs, and accepts traffic from all other VLANs on the system.

a. *<u>On which VLANs should traffic be enabled or disabled?</u>*
Use this section to specify the VLANs that should not accept client traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so it is critical in this case that you click the VLANs and then use the Move button (>>) to adjust list membership.

⚠ ***Warning*** *If you choose to disable certain VLANs, you must move at least one VLAN to the Options list. Otherwise, the system will deny traffic from all VLANs on the box, and the configuration, although valid, will not pass any traffic.*

If you are using basic mode, continue with *Finished on page 18.*

l. *<u>Which URL categories should bypass SSL filtering?</u>*
If you want specific URL categories to bypass SSL decryption, use the arrow buttons to move URL categories to or from the Selected list. Typically this is done for concerns over user privacy or for categories that contain items that may rely on specific SSL certificates to be presented as part of a verification process (e.g., software update tools).

➡ ***<u>Note:</u>*** *You must have licensed URL filtering and provisioned Secure Web Gateway (SWG) to use the URL SSL bypass feature.*

m. *<u>Do you want to apply additional iRules to decrypted SSL traffic before it is forwarded to the security device?</u>* **Advanced**
You can add custom iRules to the deployment. iRules are a scripting language that allows an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. An iRule contains the set of instructions the system uses to process data flowing through it, either in the header or payload of a packet.

Select if have preexisting iRules you want to add to your implementation.

⚠ ***Warning*** *While iRules can provide additional functionality not present in the iApp, improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommended you verify the impact of an iRule prior to deployment in a production environment.*

If you have iRules you want to attach to the virtual server the iApp creates, from the **Options** box, click the name of the applicable iRule(s) and then click the Add (**<<**) button to move them to the **Selected** box.

Advanced Firewall Manager (BIG-IP AFM)

Use this section for information about the questions in the template.

1. *<u>Do you want to use BIG-IP AFM to protect your Air Gap deployment?</u>*
If you have licensed and provisioned the BIG-IP AFM module, you can use it to help protect your air gap implementation. You can use the BIG-IP AFM to allow or deny the client IP address(es) or networks that can access external resources. Choose whether you want to deploy the BIG-IP AFM at this time.

    - **No, do not use AFM to secure your application**
    Select this option if you do not want to use BIG-IP AFM at this time, and then continue with *Finished on page 18*.

    - **Yes, use F5's recommended AFM configuration**
    Select this option if you want to deploy the BIG-IP AFM at this time, and then answer the following questions.

    a. *<u>Do you want to restrict access to external resources by client IP or network address?</u>*
    You can use the BIG-IP AFM to restrict access to external resources by either IP address or network address. If enabled, the system will only allow access to the virtual server from the address(es) you specify, and implicitly deny access to <u>all</u> unspecified addresses. You have the option to deny access to external resources in an upcoming question.

- **No, do not restrict source addresses (allow all sources)**
  Select this option if you do not want to restrict the source IP addresses or networks that can access external resources.

- **Restrict source addresses**
  Select this option if you want to restrict access to external resources. You specify the IP addresses or networks in the next question.

  a. *What client IP or network addresses should be allowed to access external resources?*
  Specify the IP or network address(es) that should be allowed external access. You can use a single IP address, a list of IP addresses separated by spaces, a range of IP addresses separated by a dash (for example 192.0.2.10-192.0.2.100), a single network address, such as 192.0.2.200/24, or any combination of these.

b. *Do you want to explicitly deny access to external resources by client IP or network address?*
You can use the BIG-IP AFM to deny access to external resources by either IP address or network address. If enabled, the system will explicitly *deny* access to the virtual server from the address(es) you specify. This is different than the previous questions in which you could specify specific addresses or networks to allow.

- **No, do not explicitly deny source addresses**
  Select this option if you do not want to deny specific source IP addresses or networks from accessing external resources.

- **Explicitly deny source addresses**
  Select this option if you want to restrict access to external resources. You specify the IP addresses or networks in the next question.

  a. *What client IP or network addresses should be explicitly denied access external resources?*
  Specify the IP or network address that should be denied external access. You can use a single IP address, a list of IP addresses separated by spaces, a range of IP addresses separated by a dash (for example 192.0.2.10-192.0.2.100), a single network address, such as 192.0.2.200/24, or any combination of these.

c. *Would you like to stage a policy for testing purposes?*
Choose whether you want to stage a firewall policy for testing purposes. A staged policy allows you to evaluate the effect a policy has on traffic by analyzing the system logs, without actually modifying traffic based on the firewall rules. You must already have a policy on the system in order to select it.

- **Do not apply a staging policy**
  Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.

- *Select an existing policy from the list*
  If you have already created a firewall policy for this implementation, select it from the list. Only policies that already exist on the system appear in the list. To create a new policy, on the Main tab, click **Security > Network Firewall > Policies**. Specific instructions for creating a firewall policy is outside the scope of this iApp and deployment guide.

d. *Which logging profile would you like to use?*
Choose whether you or not you want to use a logging profile for this AFM implementation. You can configure the BIG-IP system to log detailed information about BIG-IP system Network Firewall events and store those logs on the BIG-IP system or a remote logging server (supports syslog and Splunk). If you want to use a logging profile, we recommend creating one outside this template. The list only contains profiles with Network Firewall enabled.

- **Do not use a logging profile**
  Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.

- *Select an existing logging profile from the list*
  If you have already created a logging profile for this implementation, select it from the list. You must create a profile before it is available in the list. To create a logging profile, on the Main tab, click **Security > Event Logs > Logging Profiles**. Specific instructions for creating a logging profile is outside the scope of this iApp and deployment guide. See the online help or the *About Local Logging with the Network Firewall* chapter of the ***BIG-IP Network Firewall: Policies and Implementations*** guide for more information.

This completes the configuration for this scenario, continue with *Finished on page 18*.

- **This BIG-IP LTM will receive egress traffic from a security device**
  Select this option if the system you are configuring is external-facing and receives traffic from a security device such as a firewall.

  a. *Would you like to forward re-encrypted outbound client traffic to a pool of routers, or use the default network route?*
     Choose whether you want to forward re-encrypted traffic to a pool of routers, or if the system should use the default network route. You must have the default network route configured on your BIG-IP system before you can select it from the list.

     - *Select the default network route you created*
       If you want to send the re-encrypted traffic out the default network route, select the route from the list. Only the default network route (destination and netmask set to 0.0.0.0) appears in this list.

       ⚠ ***Warning*** *You MUST have configured this route with a gateway IP address before deploying this iApp. Selecting a default route that forwards traffic to a VLAN or pool will result in an error message.*

     - **Forward to a pool**
       Select this option if you want to forward the re-encrypted traffic to a pool of routers. You specify the router IP address(es) in the next question.

       a. *To which device(s) should this BIG-IP LTM forward re-ecrypted outbound client traffic?*
          Specify the IP address of each device to which the BIG-IP system should forward outbound client traffic it has re-encrypted. Typically, the egress device should point to an outbound router.
          Click the Add button to include more devices. The system creates load balancing pools with the addresses you specify here.

  b. *Which HTTP profile do you want to use for server-side traffic?*
     The HTTP profile contains settings for instructing the BIG-IP system how to handle HTTP traffic. Choose whether you want the iApp to create a new HTTP profile for server-side traffic or if you have previously created an HTTP profile for this deployment.

     Unless you have requirements for configuring specific HTTP settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic** >> **Profiles** : **Services : HTTP** to create a HTTP profile. To select any new profiles you create, you need to restart or reconfigure this template.

     - *Select an existing HTTP profile from the list*
       If you already created an HTTP profile for this implementation, select it from the list.

     - **Create a new HTTP profile (recommended)**
       Select this option to have the iApp to create a new HTTP profile.

  c. *Do you want to restrict server-side traffic to specific VLANs?*
     The BIG-IP system allows you to restrict client traffic to specific VLANs that are present on the system. This can provide an additional layer of security, as you can allow or deny traffic from the VLANs you choose. By default, all VLANs configured on the system are enabled. If you select to enable or disable traffic on specific VLANs, you must specify the VLANs in the next question. The VLAN objects must already be configured on this BIG-IP system before you can select them.

     - **Enable traffic on all VLANs and Tunnels**
       Choose this option to allow traffic from all VLANs and Tunnels. If you select this option, the question asking about VLANs disappears. Continue with the next question.

     - **Yes, enable traffic only on the VLANs I specify**
       Choose this option to restrict client traffic to specific VLANs that you choose in the following question. The system will accept client traffic from these VLANs, and deny traffic from all other VLANs on the system.

       a. *On which VLANs should server-side traffic be enabled or disabled?*
          Use this section to specify the VLANs that accept server-side traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so click the VLANs and then use the Move buttons to adjust list membership.

          ➡ ***Note:*** *If you choose to allow traffic from certain VLANs, when additional VLANs are added to the BIG-IP system at a later time, this iApp configuration will deny traffic from these VLANs by default. To accept traffic from these VLANs, you must re-enter the template and add the VLAN(s).*

- **Yes, disable traffic only on the VLANs I specify**
  Choose this option to deny server-side traffic from the specific VLANs that you choose in the following question. The system will refuse server-side traffic from these VLANs, and accept traffic from all other VLANs on the system.

  a. *On which VLANs should traffic be enabled or disabled?*
  Use this section to specify the VLANs that should not accept server-side traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so it is critical in this case that you click the VLANs and then use the Move button (>>) to adjust list membership.

  ⚠ ***Warning*** *If you choose to disable certain VLANs, you must move at least one VLAN to the Options list. Otherwise, the system will deny traffic from all VLANs on the box, and the configuration, although valid, will not pass any traffic.*

If you are using basic mode, continue with *Finished on page 18.*

d. *Do you want to apply additional iRules to decrypted SSL traffic before it is forwarded to the security device?* **Advanced**
You can add custom iRules to the deployment. iRules are a scripting language that allows an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. An iRule contains the set of instructions the system uses to process data flowing through it, either in the header or payload of a packet.

Select if have preexisting iRules you want to add to your implementation.

⚠ ***Warning*** *While iRules can provide additional functionality not present in the iApp, iRules are an advanced feature and should be used only if you understand how each iRule will affect your deployment, including application behavior and BIG-IP system performance.*

If you have iRules you want to attach to the virtual server the iApp creates, from the **Options** box, click the name of the applicable iRule(s) and then click the Add (**<<**) button to move them to the **Selected** box.

If you do not want to add any iRules to the configuration, continue with *Finished on page 18*.

- **This BIG-IP LTM will receive ingress and egress traffic on different networks**
  Select this option if you are using a single F5 device in this deployment, and this BIG-IP LTM will receive ingress traffic from internal clients and egress traffic from security device(s) on different networks. This is a single device solution, and you must have at least two VLANs configured on the system (one for receiving traffic from clients and one for receiving traffic from a security device).   For more information on configuring VLANs, see the Help tab or the BIG-IP documentation.

  This configuration assumes the security device is OSI layer 3-based (or higher), and that it potentially spans networks. This security device must be configured to send traffic to the internal self IP address of the egress device (in this scenario, a self IP address on a different network on the same BIG-IP device.

  a. *Which type of forward proxy are you deploying at this time?*
  Choose whether you want to deploy the system as an Explicit Forward Proxy or a Transparent Forward Proxy. Selecting Explicit Proxy configures an HTTP proxy in explicit mode. In this scenario, browser clients must be specifically configured to use the proxy via system settings. Selecting Transparent Proxy configures a transparent forwarding proxy. In this scenario, browser clients passing through the BIG-IP system do not need to be configured with proxy settings. Note that this mode requires the traffic from the clients is routed to one of the BIG-IP Self IP addresses as part of its route out of the network.

    - **Transparent Proxy**
      Select this option if you want to configure the BIG-IP system as a transparent forward proxy. Continue with *b. Which Client SSL profile (with SSL Forward proxy enabled) do you want to use?*

    - **Explicit Proxy**
      Select this option if you want to configure the system as an Explicit proxy. You must answer the following questions.

      a. *What IP address and port do you want to use for the proxy virtual server?*
      Type the IP address and port (both are required) for your explicit proxy instance. Only change the port if you have modified it from the default port, **3128**.

      b. *What is the FQDN of this proxy?*
      Type the fully qualified domain name of your explicit proxy.

      c. *Do you want the system to forward all name requests?*
      Chose whether you want the system to forward all name requests to a group of DNS servers, or if you want the BIG-IP system to resolve and cache the names directly (e.g. follow root hints). The default is to forward DNS name requests to a group of external resolvers.

        - **No, resolve all names directly**
          Choose this option if you want the BIG-IP system to resolve and cache the names directly. Continue with the next question.

        - **Yes, forward all requests**
          Choose this option if you want the system to forward all DNS requests. You must specify the servers to which you want to forward requests in the next question.

          a. *Which DNS servers do you want to use for forwarding?*
          Specify the IP address(es) of the DNS servers you will use to resolve external host names by this proxy instance. Click the Add button to include more servers. If you are using a different port than the default (53), make sure to modify the Port value.

  b. *Which Client SSL profile (with SSL Forward proxy enabled) do you want to use?*
  Specify whether you want the template to create a new Client SSL profile, or if you created a custom Client SSL profile that has SSL Forward Proxy enabled and uses a CA certificate trusted by your clients, you can select it from the list.

  ⓘ ***Important*** *The CA certificate (and private key) used here is for issuing new server certificates. The CA certificate must have the Digital Signature and Certificate Signing key usage properties (at a minimum). We recommend using a subordinate CA certificate if available.*

    - *Select the Client SSL profile you created from the list*
      If you created a custom Client SSL profile for this implementation, select it from the list. This profile must contain a certificate and key from a Certificate Authority, and they must be trusted by your internal clients. For information on creating a custom Client SSL profile and importing certificates and keys, see the Help tab or the BIG-IP documentation on *support.f5.com*.

    - **Create a new Client SSL profile**
      Select this option if you want the iApp template to create a new SSL profile as a part of this application service. Although the iApp template creates the profile, it cannot import the certificate and key. You must have already imported

a certificate and key from a Certificate Authority which are trusted by your internal clients onto this BIG-IP system. To import certificates and keys, see **System > File Management > SSL Certificate List**. For specific instructions, see the Help tab or the BIG-IP documentation on *support.f5.com*.

a. *Which trusted CA certificate do you want to use to issue server certificates for client-side connections?*
Select the CA certificate you imported onto the BIG-IP system for client authentication that is trusted by your internal clients.

⚠ *Warning* *Remember, this certificate must be issued by Certificate Authority, trusted by your internal clients., and have the Digital Signature and Certificate Signing key usage properties. We recommend using a subordinate CA certificate if available.*

b. *Which trusted CA private key do you want to use to issue server certificates for client-side connections?*
Select the SSL private key associated with the certificate you selected.

c. *Which hostnames would you like to bypass SSL interception?* Advanced
If you created a BIG-IP *Data Group* object containing host names that you want to bypass SSL interception, you can select it from the list. Only previously created Data Group objects appear in the list. Creating a Data Group is not a part of this template; if you want to create a Data Group, go to **Local Traffic > iRules > Data Group List**. For specific information, see the Help tab or the BIG-IP documentation.

- **Do not bypass hostnames**
  Select this option if you do not want any hostnames to bypass SSL interception, or if you have not yet created a Data Group with hostnames.
- *Select the Data Group list you created with the hostnames you want to bypass*
  Select the Data Group list you created for the hostnames you want to bypass.

d. *Which source IP addresses would you like to bypass SSL interception?* Advanced
If you created a BIG-IP *Data Group* object containing source IP address that you want to bypass SSL interception, you can select it from the list. Only previously created Data Group objects appear in the list. Creating a Data Group is not a part of this template; if you want to create a Data Group, go to **Local Traffic > iRules > Data Group List**. For specific information, see the Help tab or the BIG-IP documentation.

- **Do not bypass source IP**
  Select this option if you do not want any source IP addresses to bypass SSL interception, or if you have not yet created a Data Group with source IP addresses.
- *Select the Data Group list you created with the source IP addresses you want to bypass*
  Select the Data Group list you created for the source IP addresses you want to bypass.

e. *Which destination IP addresses would you like to bypass SSL interception?* Advanced
If you created a BIG-IP *Data Group* object containing destination IP addresses you want to bypass SSL interception, you can select it from the list. Only previously created Data Group objects appear in the list. Creating a Data Group is not a part of this template; if you want to create a Data Group, go to **Local Traffic > iRules > Data Group List**. For specific information, see the Help tab or the BIG-IP documentation.

- **Do not bypass destination IP addresses**
  Select this option if you do not want any destination IP addresses to bypass SSL interception, or if you have not yet created a Data Group with destination IP addresses.
- *Select the Data Group list you created with the destination IP addresses you want to bypass*
  Select the Data Group list you created for the destination IP addresses you want to bypass.

f. *Which certificate bundle contains your Trusted Root CAs?* Advanced
Select the certificate bundle that contains your Trusted Root Certificate Authorities. For this question, you can leave the default if applicable, as it contains many of the most common Certificate Authorities. You can view the list at **System > File Management > SSL Certificate List > ca-bundle**, in the **Certificate Subject(s)** field.

g. *What action should be taken for an expired certificate?* Advanced
Choose the action you want the BIG-IP system to perform if the certificate has expired.

- **Drop**
  Select this option if the BIG-IP system should drop the connection from a client with an expired certificate.
- **Ignore**
  Select this option if the BIG-IP system should ignore the expired certificate and allow the connection.

h. *What action should be taken for an untrusted certificate?* Advanced
Choose the action you want the BIG-IP system to perform if the certificate is not trusted.

- **Drop**
  Select this option if the system should drop the connection from a client with an untrusted certificate.
- **Ignore**
  Select this option if the BIG-IP system should ignore the untrusted certificate and allow the connection.

i. *To which device(s) should this BIG-IP LTM forward decrypted outbound client traffic?*
Specify the IP address of each device to which the BIG-IP system should forward outbound client traffic. For the ingress device, the destination should be the self IP address of the egress device if the security device between ingress and egress points is OSI layer 2-based. If the device is OSI layer 3-based, this could be the IP address of the security device itself, or another device configured to route traffic to the security device.

Click **Add** to include more devices. The system creates load balancing pools with the addresses you specify here.

j. *Which HTTP profile do you want to use for client-side traffic?*
The HTTP profile contains settings for instructing the BIG-IP system how to handle HTTP traffic. Choose whether you want the iApp to create a new HTTP profile or if you have previously created an HTTP profile for this deployment.

Unless you have requirements for configuring specific HTTP settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic** >> **Profiles** : **Services : HTTP** to create a HTTP profile. To select any new profiles you create, you need to restart or reconfigure this template.

- *Select an existing HTTP profile from the list*
  If you already created an HTTP profile for this implementation, select it from the list.
- **Create a new HTTP profile (recommended)**
  Select this option to have the iApp to create a new HTTP profile.

k. *Do you want to restrict client traffic to specific VLANs?*
The BIG-IP system allows you to restrict client traffic to specific VLANs that are present on the system. This can provide an additional layer of security, as you can allow or deny traffic from the VLANs you choose. By default, all VLANs configured on the system are enabled. If you select to enable or disable traffic on specific VLANs, you must specify the VLANs in the next question. The VLAN objects must already be configured on this BIG-IP system before you can select them.

- **Enable traffic on all VLANs and Tunnels**
  Choose this option to allow traffic from all VLANs and Tunnels. If you select this option, the question asking about VLANs disappears. Continue with the next question.
- **Yes, enable traffic only on the VLANs I specify**
  Choose this option to restrict client traffic to specific VLANs that you choose in the following question. The system will accept client traffic from these VLANs, and deny traffic from all other VLANs on the system.

  a. *On which VLANs should traffic be enabled or disabled?*
  Use this section to specify the VLANs that accept client traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so click the VLANs and then use the Move buttons to adjust list membership.

  ➡ ***Note:*** *If you choose to allow traffic from certain VLANs, when additional VLANs are added to the BIG-IP system at a later time, this iApp configuration will deny traffic from these VLANs by default. To accept traffic from these VLANs, you must re-enter the template and add the VLAN(s).*

- **Yes, disable traffic only on the VLANs I specify**
  Choose this option to deny client traffic from the specific VLANs that you choose in the following question. The system refuses client traffic from these VLANs, and accepts traffic from all other VLANs on the system.

  a. *On which VLANs should traffic be enabled or disabled?*
  Use this section to specify the VLANs that should not accept client traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so it is critical in this case that you click the VLANs and then use the Move button (>>) to adjust list membership.

  ⚠ ***Warning*** *If you choose to disable certain VLANs, you must move at least one VLAN to the Options list. Otherwise, the system will deny traffic from all VLANs on the box, and the configuration, although valid, will not pass any traffic.*

l. *Which URL categories should bypass SSL filtering?*
If you want specific URL categories to bypass SSL decryption, use the arrow buttons to move URL categories to or from the Selected list. Typically this is done for concerns over user privacy or for categories that contain items

that may rely on specific SSL certificates to be presented as part of a verification process (e.g., software update tools).

➡ ***Note:*** *You must have licensed URL filtering and provisioned Secure Web Gateway (SWG) to use the URL SSL bypass feature.*

m. *Do you want to apply additional iRules to decrypted SSL traffic before it is forwarded to the security device?* **Advanced**
You can add custom iRules to the deployment. iRules are a scripting language that allows an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. An iRule contains the set of instructions the system uses to process data flowing through it, either in the header or payload of a packet.

Select if have preexisting iRules you want to add to your implementation.

⚠ *Warning* *While iRules can provide additional functionality not present in the iApp, improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommended you verify the impact of an iRule prior to deployment in a production environment.*

If you have iRules you want to attach to the virtual server the iApp creates, from the **Options** box, click the name of the applicable iRule(s) and then click the Add (**<<**) button to move them to the **Selected** box.

c. *Would you like to forward re-encrypted outbound client traffic to a pool of routers, or use the default network route?*
Choose whether you want to forward re-encrypted traffic to a pool of routers, or if the system should use the default network route. You must have the default network route configured on your BIG-IP system before you can select it from the list.

- *Select the default network route you created*
If you want to send the re-encrypted traffic out the default network route, select the route from the list. Only the default network route (destination and netmask set to 0.0.0.0) appears in this list.

⚠ *Warning* *You MUST have configured this route with a gateway IP address before deploying this iApp. Selecting a default route that forwards traffic to a VLAN or pool will result in an error message.*

- **Forward to a pool**
Select this option if you want to forward the re-encrypted traffic to a pool of routers. You specify the router IP address(es) in the next question.

  a. *To which device(s) should this BIG-IP LTM forward re-ecrypted outbound client traffic?*
  Specify the IP address of each device to which the BIG-IP system should forward outbound client traffic it has re-encrypted. Typically, the egress device should point to an outbound router.
  Click the Add button to include more devices. The system creates load balancing pools with the addresses you specify here.

d. *Which HTTP profile do you want to use for server-side traffic?*
The HTTP profile contains settings for instructing the BIG-IP system how to handle HTTP traffic. Choose whether you want the iApp to create a new HTTP profile for server-side traffic or if you have previously created an HTTP profile for this deployment.

Unless you have requirements for configuring specific HTTP settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic** >> **Profiles** : **Services : HTTP** to create a HTTP profile. To select any new profiles you create, you need to restart or reconfigure this template.

- *Select an existing HTTP profile from the list*
If you already created an HTTP profile for this implementation, select it from the list.

- **Create a new HTTP profile (recommended)**
Select this option to have the iApp to create a new HTTP profile.

e. *Do you want to restrict server-side traffic to specific VLANs?*
The BIG-IP system allows you to restrict client traffic to specific VLANs that are present on the system. This can provide an additional layer of security, as you can allow or deny traffic from the VLANs you choose. By default, all VLANs configured on the system are enabled. If you select to enable or disable traffic on specific VLANs, you must specify the VLANs in the next question. The VLAN objects must already be configured on this BIG-IP system before you can select them.

- **Enable traffic on all VLANs and Tunnels**
  Choose this option to allow traffic from all VLANs and Tunnels. If you select this option, the question asking about VLANs disappears. Continue with the next question.
- **Yes, enable traffic only on the VLANs I specify**
  Choose this option to restrict client traffic to specific VLANs that you choose in the following question. The system will accept client traffic from these VLANs, and deny traffic from all other VLANs on the system.

  a. *On which VLANs should server-side traffic be enabled or disabled?*
  Use this section to specify the VLANs that accept server-side traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so click the VLANs and then use the Move buttons to adjust list membership.

  ➡ ***Note:*** *If you choose to allow traffic from certain VLANs, when additional VLANs are added to the BIG-IP system at a later time, this iApp configuration will deny traffic from these VLANs by default. To accept traffic from these VLANs, you must re-enter the template and add the VLAN(s).*

  - **Yes, disable traffic only on the VLANs I specify**
    Choose this option to deny server-side traffic from the specific VLANs that you choose in the following question. The system will refuse server-side traffic from these VLANs, and accept traffic from all other VLANs on the system.

    a. *On which VLANs should traffic be enabled or disabled?*
    Use this section to specify the VLANs that should not accept server-side traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so it is critical in this case that you click the VLANs and then use the Move button (>>) to adjust list membership.

    ⚠ ***Warning*** *If you choose to disable certain VLANs, you must move at least one VLAN to the Options list. Otherwise, the system will deny traffic from all VLANs on the box, and the configuration, although valid, will not pass any traffic.*

f. *Do you want to apply additional iRules to decrypted SSL traffic before it is forwarded to the security device?* **Advanced**
You can add custom iRules to the deployment. iRules are a scripting language that allows an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. An iRule contains the set of instructions the system uses to process data flowing through it, either in the header or payload of a packet.

Select if have preexisting iRules you want to add to your implementation.

⚠ ***Warning*** *While iRules can provide additional functionality not present in the iApp, iRules are an advanced feature and should be used only if you understand how each iRule will affect your deployment, including application behavior and BIG-IP system performance.*

If you have iRules you want to attach to the virtual server the iApp creates, from the **Options** box, click the name of the applicable iRule(s) and then click the Add (**<<**) button to move them to the **Selected** box.

Advanced Firewall Manager (BIG-IP AFM)

Use this section for information about the questions in the template.

1. ***Do you want to use BIG-IP AFM to protect your Air Gap deployment?***
If you have licensed and provisioned the BIG-IP AFM module, you can use it to help protect your air gap implementation. You can use the BIG-IP AFM to allow or deny the client IP address(es) or networks that can access external resources. Choose whether you want to deploy the BIG-IP AFM at this time.

   - **No, do not use AFM to secure your application**
     Select this option if you do not want to use BIG-IP AFM at this time, and then continue with *Finished on page 18*.
   - **Yes, use F5's recommended AFM configuration**
     Select this option if you want to deploy the BIG-IP AFM at this time, and then answer the following questions.

     a. *Do you want to restrict access to external resources by client IP or network address?*
     You can use the BIG-IP AFM to restrict access to external resources by either IP address or network address. If enabled, the system will only allow access to the virtual server from the address(es) you specify, and implicitly deny access to all unspecified addresses. You have the option to deny access to external resources in an upcoming question.

- **No, do not restrict source addresses (allow all sources)**
  Select this option if you do not want to restrict the source IP addresses or networks that can access external resources.
- **Restrict source addresses**
  Select this option if you want to restrict access to external resources. You specify the IP addresses or networks in the next question.

  a. *What client IP or network addresses should be allowed to access external resources?*
  Specify the IP or network address(es) that should be allowed external access. You can use a single IP address, a list of IP addresses separated by spaces, a range of IP addresses separated by a dash (for example 192.0.2.10-192.0.2.100), a single network address, such as 192.0.2.200/24, or any combination of these.

b. *Do you want to explicitly deny access to external resources by client IP or network address?*
You can use the BIG-IP AFM to deny access to external resources by either IP address or network address. If enabled, the system will explicitly *deny* access to the virtual server from the address(es) you specify. This is different than the previous questions in which you could specify specific addresses or networks to allow.

- **No, do not explicitly deny source addresses**
  Select this option if you do not want to deny specific source IP addresses or networks from accessing external resources.
- **Explicitly deny source addresses**
  Select this option if you want to restrict access to external resources. You specify the IP addresses or networks in the next question.

  a. *What client IP or network addresses should be explicitly denied access external resources?*
  Specify the IP or network address that should be denied external access. You can use a single IP address, a list of IP addresses separated by spaces, a range of IP addresses separated by a dash (for example 192.0.2.10-192.0.2.100), a single network address, such as 192.0.2.200/24, or any combination of these.

c. *Would you like to stage a policy for testing purposes?*
Choose whether you want to stage a firewall policy for testing purposes. A staged policy allows you to evaluate the effect a policy has on traffic by analyzing the system logs, without actually modifying traffic based on the firewall rules. You must already have a policy on the system in order to select it.

- **Do not apply a staging policy**
  Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.
- *Select an existing policy from the list*
  If you have already created a firewall policy for this implementation, select it from the list. Only policies that already exist on the system appear in the list. To create a new policy, on the Main tab, click **Security > Network Firewall > Policies**. Specific instructions for creating a firewall policy is outside the scope of this iApp and deployment guide.

d. *Which logging profile would you like to use?*
Choose whether you or not you want to use a logging profile for this AFM implementation. You can configure the BIG-IP system to log detailed information about BIG-IP system Network Firewall events and store those logs on the BIG-IP system or a remote logging server (supports syslog and Splunk). If you want to use a logging profile, we recommend creating one outside this template. The list only contains profiles with Network Firewall enabled.

- **Do not use a logging profile**
  Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.
- *Select an existing logging profile from the list*
  If you have already created a logging profile for this implementation, select it from the list. You must create a profile before it is available in the list. To create a logging profile, on the Main tab, click **Security > Event Logs > Logging Profiles**. Specific instructions for creating a logging profile is outside the scope of this iApp and deployment guide. See the online help or the *About Local Logging with the Network Firewall* chapter of the ***BIG-IP Network Firewall: Policies and Implementations*** guide for more information.

Finished

Review the answers to your questions. When you are satisfied, click the **Finished** button. The BIG-IP system creates the relevant objects.

Next Steps

After completing the iApp Template, the BIG-IP Application Services page opens for the Air Gap service you just created. To see the list of all the configuration objects created to support the implementation, on the Menu bar, click **Components**. The complete list of all related objects opens. You can click individual objects to see the settings. Once the objects have been created, you are ready to use the new deployment.

Client configuration

In a two device scenario, you must make sure that the default route of the internal clients is set to the self IP address of the internal BIG-IP system. If using a single box configuration, you must have a minimum of two self IP addresses. Clients must forward outbound traffic to the self IP address configured on the VLAN that you selected in response to the 'On which VLANs should client-side traffic be enabled or disabled?' question.

Modifying the iApp configuration

The iApp Application Service you just created can be quickly and easily modified if you find it necessary to make changes to the configuration. The Strict Updates feature of the iApp prevents users from manually modifying the iApp configuration (Strict Updates can be turned off, but use extreme caution). iApp allows you to re-enter the template, make changes, and then update the template. The modifications are automatically made to any of the associated objects.

To modify the configuration

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. Click the name of your Application Service from the list.
3. On the Menu bar, click **Reconfigure**.
4. Make the necessary modifications to the template.
5. Click the **Finished** button.

Troubleshooting

Use this section for common issues and troubleshooting steps.

➤ **SSL connection attempts fail with a bad signature error message when using OpenSSL s_client**

There is a known issue that causes SSL connection attempts to fail with an error when using **s_client**. ECDHE_ECDSA and DHE_DSS ciphers do not work with OpenSSL 1.0.1k and later.

This issue occurs when all of the following conditions are met:

- You are using BIG-IP v11.5.1, 11.5.2, or 11.6
- You have configured the BIG-IP system to process Secure Socket Layer (SSL) traffic using a Client SSL profile.
- You attempt to create a new SSL connection using OpenSSL version 1.0.1k or later **s_client** utility.
- The new SSL connections attempt to negotiate any of the ECDHE_ECDSA or DHE_DSS ciphers.

This issue has been fixed in 11.5.3 and in 11.6.0 HF5. If you are experiencing this issue, upgrade to one of those versions. For more information, see *https://support.f5.com/kb/en-us/solutions/public/16000/400/sol16461.html*.

Appendix A: Manual Configuration tables

We strongly recommend using the iApp template to configure the BIG-IP system. Users familiar with the BIG-IP system can use the following table to manually configure the BIG-IP system. The table contains a list of configuration objects along with any non-default settings you should configure as a part of this deployment. Settings not mentioned in the table can be configured as applicable for your configuration. For specific instructions on configuring individual objects, see the online help or product manuals.

➡ ***Note:*** *If you are you deploying a single box scenario where the BIG-IP LTM receives ingress and egress traffic on different networks, you must have at least two VLANs configured on the system (one for receiving traffic from clients and one for receiving traffic from a security device). To configure the system, use either the Transparent or Explicit proxy table in this section, and then the Configuration table if the BIG-IP LTM is receiving egress traffic from a security device on page 26*

Configuration table if the BIG-IP LTM is receiving ingress traffic from internal clients: Transparent Proxy

This is the Internal LTM as a transparent proxy, receiving traffic from internal clients in the two device scenario. If using a single BIG-IP LTM, this configuration should be on the VLAN with the internal clients.

Pools (*Main tab > Local Traffic > Pools*)		
Pool using the wildcard port		
Name	Type a unique name, such as **airgap-ingress-pool-wildcard**	
Load Balancing Method	**Least Connections (Member)**	
Address	Type the IP Address of a device to which the system should forward decrypted outbound client traffic	
Service Port	**0** Click **Add** to repeat Address and Port for all nodes	
Pool using port 80		
Name	Type a unique name, such as **airgap-ingress-pool-80**	
Load Balancing Method	**Least Connections (Member)**	
Address	Type the IP Address of a device to which the system should forward decrypted outbound client traffic	
Service Port	**80** Click **Add** to repeat Address and Port for all nodes	
Profiles (*Main tab > Local Traffic > Profiles*)		
HTTP *(Profiles > Services)*	Name	Type a unique name
	Parent Profile	**http**
Client SSL *(Profiles > SSL)*	Name	Type a unique name
	Parent Profile	**clientssl**
	SSL Forward Proxy	**Enabled**
	CA Certificate and Key	Select the Certificate and Key you imported from a Certificate Authority. This certificate must be trusted by your internal clients.
	Certificate Extensions List	Enable the **Extended Key Usage** extension (leave Subject Alternative Name enabled).
	SSL Forward Proxy Bypass	**Enabled**
Server SSL *(Profiles > SSL)*	Name	Type a unique name
	Parent Profile	**serverssl**
	SSL Forward Proxy	**Enabled**
	SSL Forward Proxy Bypass	**Enabled** (optional: if you are using URL filtering. Only in version 11.5 and later)
	Secure Negotiation	**Request**
	Server Certificate	**Require**
	Expire Certificate Response Control	Select **Drop** or **Ignore**
	Untrusted Certificate Response Control	Select **Drop** or **Ignore**
	Trusted Certificate Authorities	Select the certificate bundle containing your Trusted Root Certificate CAs. The default (ca-bundle.crt) contains many of the most common CAs.
iRules (*Main tab > Local Traffic > iRules*)		
iRule if you are not using SWG bypass/URL filtering		
Name	Type a unique name. **Note:** if you are using URL Filtering, do not create this iRule.	
Definition	For the iRule definition, see *iRule if you are not using SSL bypass URL filtering on page 22.*	
iRule if you are using SWG bypass/URL filtering		
Name	Type a unique name.	
Definition	For the iRule definition, see *Adding URL Filtering to your air gap solution on page 24*	

Virtual Servers (*Main tab > Local Traffic > Virtual Servers*)	
TCP virtual server	
Name	Type a unique name
Destination Address	0.0.0.0/0
Service Port	0
Protocol	TCP
HTTP Profile	Select the HTTP profile you created
SSL Profile (Client)	Select the Client SSL profile you created above
SSL Profile (Server)	Select the Server SSL profile you created
VLAN and Tunnel Traffic	Optional: You can restrict client-side traffic to specific VLANs. To use this feature, select **Enabled on** or **Disabled on**, and then move the appropriate VLANs to the **Selected** box. IMPORTANT: If you are deploying on a single device, you ***must*** select different VLANs for receiving traffic from clients and security devices.
Address Translation	Clear the check box to **Disable** Address Translation
Port Translation	Ensure the box is checked to **Enable** Port Translation.
Default Pool	Select the **Pool using the wildcard port** you created
iRules	Enable the iRule you created
UDP virtual server	
Name	Type a unique name
Destination Address	0.0.0.0/0
Service Port	0
Protocol	UDP
VLAN and Tunnel Traffic	Optional: You can restrict client-side traffic to specific VLANs. To use this feature, select **Enabled on** or **Disabled on**, and then move the appropriate VLANs to the **Selected** box. IMPORTANT: If you are deploying on a single device, you ***must*** select different VLANs for receiving traffic from clients and security devices.
Address Translation	Clear the check box to **Disable** Address Translation
Port Translation	Clear the check box to **Disable** Port Translation.
Default Pool	Select the **Pool using the wildcard port** you created

iRule if you are not using SSL bypass URL filtering

Use this iRule if you are not configuring URL filtering, or do not have an active URL Filtering subscription and have not provisioned Secure Web Gateway (SWG). Replace the pool name in line 23 with your ingress pool on port 80.

```
1   when CLIENT_ACCEPTED {
2         HTTP::disable
3         SSL::disable clientside
4         SSL::disable serverside
5         TCP::collect
6   }
7   when CLIENT_DATA {
8         binary scan [TCP::payload] c type
9         if { ( $type == 23 ) or ( $type == 20 ) } {
10            SSL::enable clientside
11            SSL::enable serverside
12        } elseif { $type == 22 } {
13            SSL::enable clientside
14            SSL::enable serverside
15            HTTP::enable
16        }
17        TCP::release
18  }
19  when HTTP_REQUEST {
20        HTTP::header insert X-Proxy-HTTPS [TCP::local_port]
21        LB::detach
22        SSL::disable serverside
23        pool <replace_with_name_of_airgap_ingress_pool_port80>
24  }
```

Configuration table if the BIG-IP LTM is receiving traffic from internal clients: Explicit Proxy

This is the Internal LTM as an Explicit receiving traffic from internal clients in the two device scenario. If using a single BIG-IP LTM, this configuration should be on the VLAN with the internal clients.

Pools (*Main tab > Local Traffic > Pools*)		
Pool using the wildcard port		
Name	Type a unique name, such as **airgap-ingress-pool-wildcard**	
Load Balancing Method	**Least Connections (Member)**	
Address	Type the IP Address of a device to which the system should forward decrypted outbound client traffic	
Service Port	**0** Click **Add** to repeat Address and Port for all nodes	
Pool using port 80		
Name	Type a unique name, such as **airgap-ingress-pool-80**	
Load Balancing Method	**Least Connections (Member)**	
Address	Type the IP Address of a device to which the system should forward decrypted outbound client traffic	
Service Port	**80** Click **Add** to repeat Address and Port for all nodes	
DNS Resolver (*Main tab > Network > DNS Resolver*)		
Name	Type a unique name	**All other settings are optional**
Tunnel (*Main tab > Network > Tunnels*)		
Name	Type a unique name	
Encapsulation Type	**tcp-forward**	**All other settings are optional**
Profiles (*Main tab > Local Traffic > Profiles*)		
HTTP: Default *(Profiles > Services)*	Name	Type a unique name
	Parent Profile	**http**
HTTP: Explicit *(Profiles > Services)*	Name	Type a unique name
	Proxy Mode	**Explicit**
	Parent Profile	**http-explicit**
	DNS Resolver	Select the DNS Resolver you created
	Tunnel Name	Select the Tunnel you created
	Host Names	Type any host names that should not be proxied
	Connection Failed Message	Optional: See *http://www.f5.com/pdf/deployment-guides/explicit-messages.zip* for our example
	DNS Lookup Failed Message	Optional: See *http://www.f5.com/pdf/deployment-guides/explicit-messages.zip* for our example
	Bad Request Message	Optional: See *http://www.f5.com/pdf/deployment-guides/explicit-messages.zip* for our example
	Bad Response Message	Optional: See *http://www.f5.com/pdf/deployment-guides/explicit-messages.zip* for our example
Server SSL *(Profiles > SSL)*	Name	Type a unique name
	Parent Profile	**serverssl**
	SSL Forward Proxy	**Enabled**
	SSL Forward Proxy Bypass	**Enabled** (optional: if you are using URL filtering. Only in version 11.5 and later)
	Secure Negotiation	**Request**
	Server Certificate	**Require**
	Expire Certificate Response Control	Select **Drop** or **Ignore**
	Untrusted Certificate Response Control	Select **Drop** or **Ignore**
	Trusted Certificate Authorities	Select the certificate bundle containing your Trusted Root Certificate CAs. The default (ca-bundle.crt) contains many of the most common CAs.
iRules (*Main tab > Local Traffic > iRules*)		
iRule if you are not using SWG bypass/URL filtering		
Name	Type a unique name. **Note:** if you are using URL Filtering, do not create this iRule.	
Definition	For the iRule definition, see *iRule if you are not using SSL bypass URL filtering on page 22.*	
iRule if you are using SWG bypass/URL filtering		
Name	Type a unique name.	
Definition	For the iRule definition, see *Adding URL Filtering to your air gap solution on page 24*	

Virtual Servers (*Main tab > Local Traffic > Virtual Servers*)	
Port 80 virtual server	
Name	Type a unique name
Destination Address	0.0.0.0/0
Service Port	80
HTTP Profile	Select the default HTTP profile you created
SSL Profile (Server)	Select the Server SSL profile you created
VLAN and Tunnel Traffic	Select **Enabled on**, and then select the Tunnel you created.
Address Translation	Clear the check box to **Disable** Address Translation
Port Translation	Ensure the box is checked to **Enable** Port Translation.
Default Pool	Select the **Pool using the wildcard port** you created
iRules	Enable the iRule you created
Port 3128 virtual server	
Name	Type a unique name
Destination Address	Type the IP address you want to use for this explicit proxy
Service Port	3128
HTTP Profile	Select the HTTP Explicit profile you created
Address Translation	Check box to **Enable** Address Translation
Port Translation	Check box to **Enable** Port Translation.
Wildcard virtual server	
Name	Type a unique name
Type	**Performance (Layer 4)**
Destination Address	0.0.0.0/0
Service Port	0
Protocol Profile (client)	Select the default **fastl4** profile, or a custom fastl4 profile if you have created one
VLAN and Tunnel Traffic	Select **Enabled on**, and then select the Tunnel you created.
Address Translation	Clear the check box to **Disable** Address Translation
Port Translation	Clear the check box to **Disable** Port Translation.
Default Pool	Select the **Pool using the wildcard port** you created

Adding URL Filtering to your air gap solution

You can optionally add URL filtering to the solution, if you have an active URL Filtering subscription and have provisioned Secure Web Gateway (SWG). To configure URL Filtering, use the following iRule. You should modify the SSL bypass categories in lines 3-7 to add or remove appropriate categories. For a full list of available categories, see **Access Policy > Secure Web Gateway > URL Categories** for a complete list of available options. You must also replace the name of the ingress pool on port 80 in line 99.

```
1   when RULE_INIT {
2      set static::ssl_bypass_categories {
3         /Common/Financial_Data_and_Services
4         /Common/Hosted_Business_Applications
5         /Common/Information_Technology
6         /Common/Online_Brokerage_and_Trading
7         /Common/Abortion
8         # you can add or remove Categories. See Access Policy - Secure Web Gateway - URL Categories for a list of options
9      }
10  }
11  when CLIENT_ACCEPTED {
12        HTTP::disable
13        SSL::disable clientside
14        SSL::disable serverside
15        TCP::collect
16  }
```

Important: This iRule continues on the next page

```
when CLIENT_DATA {
      binary scan [TCP::payload] c type
      if { ( $type == 23 ) or ( $type == 20 ) } {
         SSL::enable clientside
         SSL::enable serverside
      } elseif { $type == 22 } {
         SSL::enable clientside
         SSL::enable serverside
         HTTP::enable
      }
      TCP::release
}
when CLIENTSSL_CLIENTHELLO {
    set sni_exists [SSL::extensions exists -type 0]
    if { $sni_exists } {
        binary scan [SSL::extensions -type 0] S1S1S1cS1a* ssl_ext_type ssl_ext_len ssl_ext_sn_list_len ssl_ext_sn_type ssl_ext_sn_len ssl_ext_sn
    }
}
when SERVERSSL_HANDSHAKE {
   if { not $sni_exists } {
      set ssl_bypass_mitm 0
      set subject [X509::subject [SSL::cert 0]]
      regexp {CN=(.*?),} $subject fullcn subcn
      if { [info exists subcn] } {
         set this_uri "http://$subcn/"
         set reply [getfield [CATEGORY::lookup $this_uri] " " 1]
         set decision [lsearch -exact $static::airgap_ssl_bypass_categories $reply]
         if {[lsearch -exact $static::airgap_ssl_bypass_categories $reply] >= 0}{
            set ssl_bypass_mitm 1
         } else {
            set ssl_bypass_mitm 0
         }
      } else {
         regexp {CN=(.*?)$} $subject fullcn subcn
         if { [info exists subcn] } {
            set this_uri "http://$subcn/"
            set reply [getfield [CATEGORY::lookup $this_uri] " " 1]
            set decision [lsearch -exact $static::airgap_ssl_bypass_categories $reply]
            if {[lsearch -exact $static::airgap_ssl_bypass_categories $reply] >= 0}{
               set ssl_bypass_mitm 1
            } else {
               set ssl_bypass_mitm 0
            }
         } else {
            set ssl_bypass_mitm 0
         }
      }
   }
}
when CLIENTSSL_SERVERHELLO_SEND {
    if { not [info exists ssl_bypass_mitm] && [info exists ssl_ext_sn] } {
        set this_uri "http://$ssl_ext_sn/"
        set reply [getfield [CATEGORY::lookup $this_uri] " " 1]
        set decision [lsearch -exact $static::airgap_ssl_bypass_categories $reply]
        if {[lsearch -exact $static::airgap_ssl_bypass_categories $reply] >= 0}{
            set ssl_bypass_mitm 1
        } else {
            set ssl_bypass_mitm 0
        }
    }

    if { [info exists ssl_bypass_mitm] } {
        if { $ssl_bypass_mitm } {
         SSL::forward_proxy policy bypass
         catch { HTTP::disable }
        } else {
         SSL::forward_proxy policy intercept
        }
    } else {
   }
}
when SERVER_CONNECTED {
    if { [info exists ssl_bypass_mitm] } {
        if { $ssl_bypass_mitm } {
         catch { HTTP::disable }
        }
    }
}
when HTTP_REQUEST {
      HTTP::header insert X-Proxy-HTTPS [TCP::local_port]
      LB::detach
      SSL::disable serverside
      pool <your_airgap_ingress_pool_port_80>
}
```

Configuration table if the BIG-IP LTM is receiving egress traffic from a security device

This is the external LTM receiving traffic from the security device(s) in the two device scenario. If using a single BIG-IP LTM, this configuration should be on the VLAN with the security devices.

Pools (*Main tab > Local Traffic > Pools*)

Pool using the wildcard port	
Name	Type a unique name, such as **airgap-egress-pool-any**
Load Balancing Method	**Least Connections (Member)**
Address	Type the IP Address of a device to which the system should forward re-encrypted outbound client traffic
Service Port	**0** Click **Add** to repeat Address and Port for all nodes

Profiles (*Main tab > Local Traffic > Profiles*)

HTTP (*Profiles > Services*)	Name Parent Profile	Type a unique name **http**
TCP (*Profiles > Protocol*)	Name Parent Profile	Type a unique name **tcp**
OneConnect (*Profiles > Protocol*)	Name Parent Profile	Type a unique name **onconnect**
Server SSL (*Profiles > SSL*)	Name Parent Profile	Type a unique name **serverssl-insecure-compatible**

iRules (*Main tab > Local Traffic > iRules*)

Create __one__ of the following iRules, depending on whether you want to forward re-encrypted outbound client traffic to a pool of routers, or use the default network route. If you use the default network route, it must already exist on the system. See the Help tab or the documentation for specific information on routes.

iRule if you want to forward re-encrypted outbound client traffic to a pool of routers	
Name	Type a unique name. Replace the text in red with the name of the pool using the wildcard port you created.
Definition	

```
when HTTP_REQUEST {
   if { not ( [HTTP::header exists X-Proxy-HTTPS] ) } {
          SSL::disable serverside
          pool <name_of_your_airgap_egress_pool_any>
   } else {
          node [lindex [active_nodes -list airgap_egress_pool_any] 0] [HTTP::header X-Proxy-HTTPS]
          HTTP::header remove X-Proxy-HTTPS
   }
}
```

iRule if you want to forward re-encrypted outbound client traffic to the default network route on the BIG-IP system	
Name	Type a unique name. Replace the text in red with the default network route on your BIG-IP system.
Definition	

```
when HTTP_REQUEST {
   if { not ( [HTTP::header exists X-Proxy-HTTPS] ) } {
          SSL::disable serverside
   } else {
          node <IP-address-of-your-default-network-route> [HTTP::header X-Proxy-HTTPS]
          HTTP::header remove X-Proxy-HTTPS
   }
}
```

Virtual Servers (*Main tab > Local Traffic > Virtual Servers*)

Port 80 virtual server	
Name	Type a unique name
Destination Address	**0.0.0.0/0**
Service Port	**80**
Protocol Profile (client)	Select the TCP profile you created
HTTP Profile	Select the HTTP profile you created
SSL Profile (Server)	Select the Server SSL profile you created
VLAN and Tunnel Traffic	You can restrict client-side traffic to specific VLANs. To use this feature, select **Enabled on** or **Disabled on**, and then move the appropriate VLANs to the **Selected** box. __IMPORTANT:__ If you are deploying on a single device, you ***must*** select different VLANs for receiving traffic from clients and security devices.
Default Pool	**None**. Do __not__ select a pool for this virtual server.
iRules	Enable the iRule you created

Wildcard virtual server	
Name	Type a unique name
Type	**Performance (Layer 4)**
Destination Address	**0.0.0.0/0**
Service Port	**0**
Protocol Profile (client)	Select the default **fastl4** profile, or a custom fastl4 profile if you have created one
VLAN and Tunnel Traffic	You can restrict client-side traffic to specific VLANs. To use this feature, select **Enabled on** or **Disabled on**, and then move the appropriate VLANs to the **Selected** box. **IMPORTANT:** If you are deploying on a single device, you ***must*** select different VLANs for receiving traffic from clients and security devices.
Default Pool	Select the **Pool using the wildcard port** you created

This completes the LTM configuration.

Manually configuring the BIG-IP AFM

If you are deploying the BIG-IP system to receive traffic from internal clients (the traffic ingress point coming from the client), or if the LTM is receiving traffic from both internal clients and a security device, you can use the BIG-IP AFM module to help protect the implementation. The AFM module must be fully licensed and provisioned before attempting this configuration.

BIG-IP AFM is particularly useful if you want to only allow access from specific clients or networks. Because this configuration can be complex, we recommend using the iApp template in version 11.6 and later to configure BIG-IP AFM.

Network Firewall settings

When configuring the BIG-IP Advanced Firewall Manager, you may want to configure your BIG-IP system to drop all traffic that you have not specifically allowed with firewall rules. This in known as ***firewall mode***. By default, your BIG-IP system is set to default-accept, or ***ADC mode***. Instructions for configuring your BIG-IP system, and the implications to consider, can be found on AskF5. For example, for BIG-IP v11.5: *http://support.f5.com/kb/en-us/products/big-ip-afm/manuals/product/network-firewall-policies-implementations-11-5-0/1.html*

If you have licensed IP Intelligence on the BIG-IP system, you can prohibit connections from sources with low reputation scores.

The following instructions cover a basic firewall configuration that is effective for the most common scenario of wanting to allow connections from a single trusted network. If you have complex requirements, such as the need to schedule different policies for different times of the day, or you want to create complicated rule or address lists, consult the BIG-IP AFM documentation. The basic steps for Policy and Rule creation apply to all scenarios.

To configure the BIG-IP AFM to allow connections from a single trusted network

1. Create a Network Firewall Policy:
    a. From the Configuration utility, click **Security > Network Firewall > Policies**, and then click **Create**.
    b. In the **Name** field, type a unique name for the policy.
    c. Click **Finished**.

2. Create a rule to allow authorized hosts or networks to connect:
    a. Click **Security > Network Firewall > Policies**.
    b. Click the name of the policy you just created.
    c. In the Rule section (below the General Properties section), click the **Add** button.
    d. Leave the **Type** list set to Rule.
    e. From the **Order** list, select **First**. The Order list only appears in version 11.5 and later. In 11.4.x, you must reorder the rules from the Policy General Properties page.
    f. In the **Name** field, type a unique name, for instance **airgap-traffic-allowed**.
    g. Ensure the **State** list is set to **Enabled**.
    h. From the **Protocol** list, select **TCP**. Leave the box to the right of TCP set to **6**.
    i. In the **Source** section, from the **Address/Region** list, select **Specify**.
    You are now able to list the trusted source addresses for your connection.
    In the following example, we will configure a single subnet as trusted.
        - Select **Address**.
        - In the box, type the network address you want to allow, including netmask if more than a single host. Specify a network using CIDR notation, such as **10.0.0.0/24**.
        - Do not configure a source port.
        - Optional: If you want to limit inbound connections to a specific VLAN or Tunnel, from the **VLAN / Tunnel** list, select **Specify**, and then move the VLANs or tunnels that are allowed access to the Selected box.
        - Click **Add**.
        - Repeat these steps for additional hosts or networks. Use **Address List** or **Address Range** when appropriate.

j. In the **Destination** section, leave the **Address/Region** and **Port** set to **Any**. Because you will be applying your policy to a virtual server that listens only on a single desired address and port, do not specify that information here.

k. If necessary, from the **Action** list, select **Accept**.

l. *Optional:* If you have configured a logging profile and want to log connections, from the **Logging** list, select **Enabled**. Typically, allowed connections do not need to be logged.

m. Click **Finished**.

3. Creating a firewall rule to block all other traffic
The next task is to create a firewall rule to block all other traffic that you have not allowed. Although this is not a required step if your BIG-IP system is set to default deny (**Firewall mode**), it is required in default-accept (**ADC mode**), and is a good practice to always configure such a rule.

a. Click **Security > Network Firewall > Policies**.

b. Click the name of the policy you created in step 1.

c. In the Rule section (below the General Properties section), click the **Add** button.

d. Leave the **Type** list set to **Rule**.

e. Leave the **Order** list, select **Last**.

f. In the **Name** field, type a unique name, for example **airgap-traffic-prohibited**.

g. Ensure the **State** list is set to **Enabled**.

h. From the **Protocol** list, select **TCP**. Leave the box to the right of TCP set to **6**.

i. In the **Source** section, leave all the lists set to **Any**

j. From the **Action** list, select either **Drop** (to silently discard incoming connections) or **Reject** (to send a Destination Unreachable message to the sender).

k. If you configured a logging profile as described in *Optional: Configuring the BIG-IP system to log network firewall events on page 30*, from the **Logging** list, select **Enabled**.

l. Click **Finished**. You return to the Policy Properties page.

m. On the Policy Properties page, in the Rules section, ensure the rule with the Action of Accept comes before the Drop or Reject rule you just created. If it does not, use the **Reorder** button and drag the rules into the correct order.

4. Apply Your Firewall Policy to your Virtual Server

a. Click **Security > Network Firewall > Active Rules**.

b. In the Rule section (below the General Properties section), click the **Add** button.

c. From the **Context** list, select **Virtual Server**, and then select the virtual server you created. If you using the Explicit proxy mode, this is the virtual server on port 3128. If you are using Transparent proxy mode, you select both the virtual servers you created that are receiving traffic from clients.

d. From the **Type** list, select **Policy**, and then select the firewall policy you created.

e. From the **Policy Type** list, select **Enforced**.

f. Click **Finished**.

Optional: Assigning an IP Intelligence Policy to your Air Gap deployment

If you want to restrict access to your air gap implementation based on the reputation of the remote sender, you can enable and assign an IP Intelligence policy. This requires an IP intelligence license; contact your F5 Sales representative for more information.

It is outside the scope of this document to provide instructions on configuring an IP Intelligence Policy. Full documentation on enabling and configuring the IP Intelligence feature can be found on AskF5. For example, the manual for BIG-IP AFM v11.5 is: *https://support.f5.com/kb/en-us/products/big-ip-afm/manuals/product/network-firewall-policies-implementations-11-5-0/5.html*

After you have enabled and configured an IP Intelligence policy, use the following steps to assign the policy to your virtual server.

To assign the IP intelligence policy to the virtual server

1. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.
2. Click the name of the applicable virtual server you created. If you using the Explicit proxy mode, this is the virtual server on port 3128. If you are using Transparent proxy mode, you select both the virtual servers you created that are receiving traffic from clients.
3. From the **Security** menu, choose **Policies**.
4. Next to **IP Intelligence**, select **Enabled**, then select the IP intelligence policy to apply to traffic on the virtual server.
5. Click **Update**. The list screen and the updated item are displayed. The IP Intelligence policy is applied to traffic on the virtual server.

Optional: Configuring the BIG-IP system to log network firewall events

If you are using BIG-IP AFM, you have the option of logging network firewall events to one or more remote syslog servers (recommended) or to log events locally. You can either use an iApp template to create the logging profile, or create the logging profile manually.

For specific information on logging on the BIG-IP system, see the appropriate guide for your version. For example, for 11.5.0:

- Remote High-Speed Logging: *https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos-implementations-11-5-0/22.html*
- Local logging: *https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos-concepts-11-5-0/11.html*

Creating the logging profile using the iApp template

Use this section to create the logging profile using the logging profile iApp template. If you have not already downloaded the iApp template, see *https://devcentral.f5.com/wiki/iApp.F5-Remote-Logging-iApp.ashx*.

To configure the logging profile iApp

1. Log on to the BIG-IP system.
2. On the Main tab, click **iApp > Application Services**.
3. Click **Create**. The Template Selection page opens.
4. In the **Name** box, type a name. In our example, we use **logging-iapp_.**
5. From the **Template** list, select **f5.remote_logging.v*<latest-version>***. The template opens
6. Use the following table for guidance on configuring the iApp template. Questions not mentioned in the table can be configured as applicable for your implementation.

Question	Your selection
Do you want to create a new pool of remote logging servers, or use an existing one?	Unless you have already created a pool on the BIG-IP system for your remote logging servers, select **Create a new pool**.
Which servers should be included in this pool?	Specify the IP addresses of your logging servers. Click **Add** to include more servers.
What port do the pool members use?	Specify the port used by your logging servers, typically **514**.
Do the pool members expect UDP or TCP connections?	**TCP**
Do you want to create a new monitor for this pool, or use an existing one?	Unless you have already created a health monitor for your pool of logging servers, select **Use a simple ICMP (ping) monitor**.
Do your log pool members require a specific log format?	If your logging servers require a specific format, select the appropriate format from the list.

7. Click **Finished**.
8. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.

9. Click the name of the applicable virtual server you created. If you using the Explicit proxy mode, this is the virtual server on port 3128. If you are using Transparent proxy mode, you select both the virtual servers you created that are receiving traffic from clients.
10. From the **Security** menu, choose **Policies**.
11. Next to **Log Profile**, select **Enabled**, then select the Logging profile you created.
12. Click **Update**. The list screen and the updated item are displayed.

➡ ***Note:*** *The iApp template creates a log publisher and attaches it to the logging profile. If the publisher does not appear in the BIG-IP Configuration utility (GUI), you can verify the configuration by running the following command from the Traffic Management shell (tmsh):* `list security log profile <your profile name>`.

Creating logging profile manually

If you do not want to use the iApp template to create a logging profile, use this section for guidance on configuring the logging profile manually. You must have access to the tmsh command line to use this method.

To manually configure a logging profile

1. Use the following guidance for configuring a health monitor and load balancing pool for the logging servers.

BIG-IP LTM Object	Non-default settings/Notes	
Health Monitor (*Local Traffic -->Monitors*)	***Name***	Type a unique name
	Type	**ICMP**
	Interval	**30** (recommended)
	Timeout	**91** (recommended)
Pool (*Local Traffic -->Pools*)	***Name***	Type a unique name
	Health Monitor	Select the appropriate monitor you created
	Slow Ramp Time	**300**
	Load Balancing Method	Choose a load balancing method. We recommend **Least Connections (Member)**
	Address	Type the IP Address of a server.
	Service Port	Type the appropriate port, such as UDP port **514**, the port on which logging typically occurs. Click **Add**, and then repeat Address and Port for all nodes

2. Log into the BIG-IP system using the command line. Enter the tmsh shell, by typing **tmsh** from the prompt.
3. Create a Remote High Speed Log (HSL) destination:

```
(tmos)# create / sys log-config destination remote-high-speed-log [name] pool-name [specified pool] protocol [udp or tcp]
```

4. If you have a specific log format requirement, create a format-specific log destination, and forward that to the previously-created HSL destination:

```
(tmos)# create / sys log-config destination [splunk|arcsight|remote-high-speed-log] [name] forward-to [HSL name]
```

5. Create a log publisher:

```
(tmos)# create / sys log-config publisher [name] destinations add { [logdestination name] }
```

6. Create the logging profile to tie everything together.
   If you chose to log allowed connections, include the green text (as in step 2 substep l in *To configure the BIG-IP AFM to allow connections from a single trusted network on page 28)*.
   If you set the rule to drop incoming connections, include the text in blue.
   If you chose to log IP intelligence events, include the text in red to add the parameter that sets the log publisher.

```
(tmos)# create / security log profile [name] network add { [name] { filter { log-acl-match-accept enabled log-acl-match-drop enabled log-acl-match-reject enabled } format { field-list { date_time action drop_reason protocol src_ip src_port dest_ip dest_port } type field-list } publisher [logpublisher name] } } ip-intelligence { log-publisher [logpublisher name] }
```

Assigning the logging profile to the virtual server

The final task is to assign the logging profile to the virtual server.

To assign the logging profile to the virtual server

1. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.
2. Click the name of the applicable virtual server.
3. From the **Security** menu, choose **Policies**.
4. Next to **Log Profile**, select **Enabled**, then select the Logging profile you created.
5. Click **Update**. The list screen and the updated item are displayed.

Document Revision History

Version	Description	Date
1.0	New guide	01-15-2015
1.1	Updated this guide for iApp v1.0.0.rc2. This update includes the ability to select the type of proxy (Transparent or Explicit), and the option to use BIG-IP AFM to protect the implementation.	03-06-2015
1.2	Updated this guide for iApp f5.airgap_egress.v1.0.0.rc4. This update includes the ability to use a default network gateway for egress traffic, and selecting pre-existing data groups for bypassing SSL intercept. The manual configuration has been updated to reflect these changes, including updated iRules.	07-06-2015
1.3	Added *Troubleshooting on page 20*, with an entry regarding SSL connection failures when using OpenSSL s_client.	07-29-2015

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

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com