



Deploying the BIG-IP System v10 with Oracle's BEA WebLogic

Version 1.0

Table of Contents

Deploying the BIG-IP system v10 with Oracle's BEA WebLogic

Prerequisites and configuration notes	1-1
Product versions and revision history	1-2
Configuration example	1-3
Configuring the BIG-IP system for BEA WebLogic	1-4
Running the BEA WebLogic application template	1-4
SSL Certificates on the BIG-IP system	1-9

Manually configuring the BIG-IP LTM for WebLogic

Creating the HTTP health monitor	2-1
Creating the pool	2-2
Creating profiles	2-4
Creating the virtual server	2-7
Manually configuring the BIG-IP LTM to offload SSL	2-10
Using SSL certificates and keys	2-10
Creating a Client SSL profile	2-11
Creating a new HTTP profile	2-11

Manually configuring the F5 WebAccelerator module with WebLogic3-1

Prerequisites and configuration notes	3-1
Configuration example	3-1
Configuring the WebAccelerator module	3-2
Creating an HTTP Class profile	3-2
Modifying the Virtual Server to use the Class profile	3-3
Creating an Application	3-4



I

Deploying the BIG-IP System v10 with BEA WebLogic

- Configuring the BIG-IP system for BEA WebLogic
- Running the BEA WebLogic application template
- SSL Certificates on the BIG-IP system

Deploying the BIG-IP system v10 with Oracle's BEA WebLogic

Welcome to the F5 and Oracle BEA® WebLogic® Server deployment guide. F5 provides a highly effective way to optimize and direct traffic for WebLogic Server with the BIG-IP Local Traffic Manager (LTM) and WebAccelerator.

BEA WebLogic Server is at the core of today's most reliable enterprise applications. F5 provides a secure, highly available and scalable application delivery networking device. This strong interoperability and integration provides a solution that delivers unparalleled traffic management functionality for those deploying services and applications on the WebLogic Enterprise Platform™.

New in version 10.0 of the BIG-IP system are Application Ready Templates. These application templates ease the process of configuring the BIG-IP system. Instead of having to individually create each object that pertains to the type of application traffic you want the BIG-IP system to manage, you can run an application template. The application template automatically creates BIG-IP system objects that are customized for that application. These objects can be either local traffic objects, TMOS objects, or both.

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

Prerequisites and configuration notes

All of the procedures in this Deployment Guide are performed on the BIG-IP system. The following are prerequisites for this solution:

- ◆ For this Deployment Guide, the BEA WebLogic Server should be running version 5.1 or 8.1.
- ◆ For this deployment guide, the BIG-IP LTM system must be running version 10.0 or later. If you are using a previous version of the BIG-IP LTM system see the *Deployment Guide* index.
- ◆ If you are using the BIG-IP system to offload SSL, we assume you have already obtained an SSL certificate and key, but it is not yet installed on the BIG-IP LTM system. For more information, see *SSL Certificates on the BIG-IP system*, on page 9.
- ◆ While we strongly recommend using the application template, you can also manually configure the BIG-IP system. For more information, see *Manually configuring the BIG-IP LTM for WebLogic*, on page 2-1.

◆ Important

All local traffic objects that an application template creates reside in administrative partition Common. Consequently, to use the application templates feature, including viewing the Templates list screen, you must have a user role assigned to your user account that allows you to view and manage objects in partition Common.

Product versions and revision history

Product and versions tested for this deployment guide:

Product Tested	Version Tested
BIG-IP System (LTM and WebAccelerator)	10.0
BEA WebLogic	5.1 and 8.1

Revision history:

Document Version	Description
1.0	New deployment guide

Configuration example

Using the configuration in this guide, the BIG-IP LTM system is optimally configured to load balance traffic to BEA WebLogic servers. Figure 1.1 shows a typical configuration with a redundant pair of BIG-IP devices with the WebAccelerator module, a cluster of WebLogic servers, and a WebLogic administration node

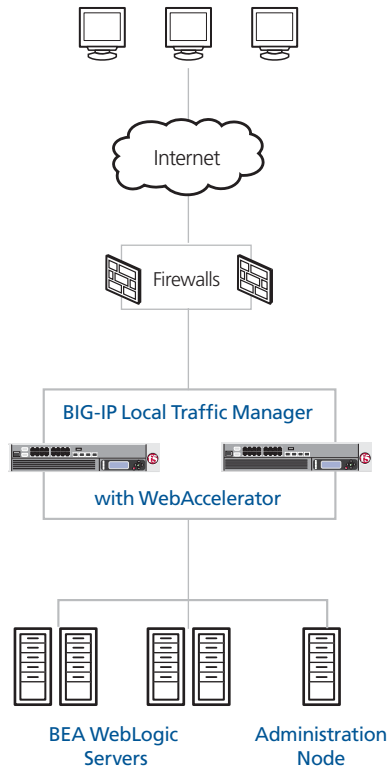


Figure 1.1 Logical configuration example

Configuring the BIG-IP system for BEA WebLogic

You can use the new Application Template feature on the BIG-IP system, to efficiently configure a set of objects corresponding to BEA WebLogic. The template uses a set of wizard-like screens that query for information and then creates the required objects. For example, depending on the settings you specify, this template creates two virtual servers, one HTTPS profile, two TCP profiles, one Persistence profile, one Client SSL profile, a OneConnect profile, one iRule, one pool, and one HTTP monitor. At the end of the template configuration process, the system presents a list of the objects created and a description for how each object interacts with the application.

◆ **Note**

Depending on which modules are licensed on your BIG-IP system, some of the options in the template may not appear.

Running the BEA WebLogic application template

To run the BEA WebLogic application template, use the following procedure. For more information on specific settings, see the online help.

To run the BEA WebLogic application template

1. Verify that your current administrative partition is set to **Common**. The Partition list is in the upper right corner.
2. On the Main tab, expand **Templates and Wizards**, and then click **Templates**. The Templates screen opens, displaying a list of templates.
3. In the Application column, click **BEA WebLogic**. The BEA WebLogic application template opens.
4. In the Virtual Server Questions section, complete the following:
 - a) You can type a unique prefix for your BEA WebLogic objects that the template will create. In our example, we leave this setting at the default, **my_weblogic**.
 - b) Enter the IP address for this virtual server. The system creates a virtual server named **<prefix from step a>_virtual_server**. In our example, we type **192.168.11.100**.
 - c) If the servers can communicate with the clients using a route through the BIG-IP system to deliver response data to the client, select **Yes** from the list. In this case, the BIG-IP does **not** translate the client's source address.

If the BIG-IP system should translate the client's source address to an address configured on the BIG-IP system, leave the list at the default setting, **No**. Selecting **No** means the BIG-IP system

will use SNAT automap. See the Online Help for more information.

In our example, we leave this at the default setting: **No**.

5. In the SSL Offload section, complete the following
 - a) if you are not using the BIG-IP system to offload SSL, leave this setting at the default, **No**. Continue with Step 6.

If you are using the BIG-IP system to offload SSL from the WebLogic devices, select **Yes** from the list. The SSL options appear.

- b) From the **Certificate** list, select the appropriate certificate you want to use for this deployment. If you plan to use a third party certificate, but have not yet installed it on the BIG-IP system, see *SSL Certificates on the BIG-IP system*, on page 1-9.
 - c) From the **Key** list, select the appropriate key for the certificate. If you have not yet installed the key on the BIG-IP system, see *SSL Certificates on the BIG-IP system*, on page 1-9.

For information on generating certificates, or using the BIG-IP LTM to generate a request for a new certificate and key from a certificate authority, see the **Managing SSL Traffic** chapter in the *Configuration Guide for Local Traffic Management*.

The screenshot shows a configuration wizard for the BEA WebLogic Template. The breadcrumb path is 'Templates and Wizards >> Templates >> BEA WebLogic'. The main title is 'BEA WebLogic Template'. A welcome message states: 'Welcome to the BEA WebLogic Template. This wizard will create a complete configuration optimized for managing BEA WebLogic traffic.'

The 'Virtual Server Questions' section contains three items:

- Unique prefix name for all objects that will be created by this template? (Text input: my_weblogic)
- What IP Address do you want to use for this BEA WebLogic virtual server? (Text input: 192.168.11.100)
- Do the BEA WebLogic servers have a route back to application clients via this BIG-IP system? (Dropdown menu: No)

The 'SSL Offload Questions' section contains three items:

- Do you want the BIG-IP system to offload SSL from the BEA WebLogic servers? (Dropdown menu: Yes)
- Certificate to authenticate the server? (You may need to import a certificate before deploying this Template.) (Dropdown menu: weblogic-ssl)
- Key used for encryption? (You may need to import a key before deploying this Template.) (Dropdown menu: weblogic-ssl)

Figure 1.2 Configuring the BIG-IP system for SSL Offload

6. In the Load Balancing Questions section, complete the following:
 - a) From the Load Balancing Method list, select an appropriate load balancing method. In our example, we leave this setting at the default, **Least Connections (member)**.
 - b) Next, add each of the BEA WebLogic devices that are a part of this deployment.

In the **Address** box, type the IP address of the first WebLogic device. In our example, we type **10.132.81.100**.

In the **Service Port** box, leave the port at **7001** (the default port for WebLogic) unless you have modified the configuration on your WebLogic servers.

Click the **Add** button. Repeat this step for each of the WebLogic devices.
 - c) Next, type a number of seconds that the BIG-IP system issues the health check. In our example, we leave this at the default level, **30**.
 - d) If you have a specific HTTP request you would like to add to the health check, type it in the box after **GET /**. This is optional.
 - e) Select the HTTP version that the BEA WebLogic servers expect clients to use. In our example, we select **Version 1.1**.

A new row appears asking for the fully qualified DNS name (FQDN) that clients use to access the WebLogic devices. In the box, type the FQDN for your WebLogic deployment. Note that this FQDN should resolve to the virtual server on the BIG-IP system. In our example, we type **weblogic.siterequest.com**.

- f) If you entered an HTTP request in step d, and want to enter a response string, type it here. This is optional. See Figure 1.3.

Load Balancing Questions	
Which load balancing method would you like to use?	Least Connections (member)
Please add the servers that will comprise this virtual server (the virtual will not be available until at least one server is added):	Address: 10.132.81.103
	Service Port: 7001 Select...
	Add
	R:1 P:1 10.132.81.100:7001 R:1 P:1 10.132.81.101:7001 R:1 P:1 10.132.81.102:7001 R:1 P:1 10.132.81.103:7001
	Edit Delete
How often should each BEA WebLogic server's health be checked?	30 seconds
HTTP request that should be sent to check server health? (HTTP 1.1 headers will be automatically added.)	GET /
What HTTP version do your BEA WebLogic servers expect clients to use?	Version 1.1
Fully qualified DNS name HTTP 1.1 clients are expected to use to access the BEA WebLogic?	weblogic.siterequ
String that should be contained within the health check response for the server to be considered healthy?	

Figure 1.3 Configuring the Load Balancing options

7. In the Protocol and Security Questions section, complete the following
 - a) If most clients will be connecting to the virtual server from a WAN, select **WAN** from the list. If most clients will be connecting from a LAN, select **LAN** from the list. This option determines the profile settings that control the behavior of a particular type of network traffic, such as HTTP connections.
 - b) If you want to use the WebAccelerator module to accelerate the WebLogic traffic, select **Yes** from the list. If you do not want to use the WebAccelerator, select **No**. This option does not appear if you do not have the WebAccelerator module licensed. The WebAccelerator module can significantly improve performance for WebLogic deployments.
 - c) If you are using the WebAccelerator module, in the **Host** box, type the fully qualified DNS name (FQDN) that your users will use to access the WebLogic deployment (the WebAccelerator application object's Requested Hosts field). Click the **Add**

button. If you have additional host names, type each one in the **Host** box, followed by clicking the **Add** button. In our example, we type **weblogic.siterequest.com** and click the **Add** button.

8. Click the **Finished** button.

Protocol Optimization and Security Questions

Will clients be connecting to this virtual server primarily over a LAN or a WAN? WAN

Would you like to use the Web Accelerator module to accelerate your BEA WebLogic traffic? Yes

Please enter the fully qualified DNS names your end users will use to access the BEA WebLogic Virtual Server (e.g., weblogic.f5.com).

Host: weblogic.siterequest.com

Add

weblogic.siterequest.com

Delete

Cancel Finished

Figure 1.4 Configuring the optimization options

After clicking **Finished**, the BIG-IP system creates the relevant objects. You see a summary screen that contains a list of all the objects that were created. This completes the BIG-IP system configuration for BEA WebLogic.

SSL Certificates on the BIG-IP system

Before you can enable the BIG-IP LTM system to act as an SSL proxy, you must install a SSL certificate on the virtual server that you wish to use for WebLogic connections on the BIG-IP LTM device. For this Deployment Guide, we assume that you already have obtained an SSL certificate, but it is not yet installed on the BIG-IP LTM system. For information on generating certificates, or using the BIG-IP LTM to generate a request for a new certificate and key from a certificate authority, see the **Managing SSL Traffic** chapter in the *Configuration Guide for Local Traffic Management*.

Importing keys and certificates

Once you have obtained a certificate, you can import this certificate into the BIG-IP LTM system using the Configuration utility. By importing a certificate or archive into the Configuration utility, you ease the task of managing that certificate or archive. You can use the Import SSL Certificates and Keys screen only when the certificate you are importing is in Privacy Enhanced Mail (PEM) format.

To import a key or certificate

1. On the Main tab, expand **Local Traffic**.
2. Click **SSL Certificates**. The list of existing certificates displays.
3. In the upper right corner of the screen, click **Import**.
4. From the **Import Type** list, select the type of import (Certificate or Key).
5. In the **Certificate (or Key) Name** box, type a unique name for the certificate or key.
6. In the **Certificate (or Key) Source** box, choose to either upload the file or paste the text.
7. Click **Import**.

If you imported the certificate, repeat this procedure for the key.



2

Manually Configuring the BIG-IP LTM System for BEA WebLogic

- Creating the HTTP health monitor
- Creating the pool
- Creating profiles
- Creating the virtual server
- Using SSL certificates and keys
- Creating a Client SSL profile
- Creating a new HTTP profile

Manually configuring the BIG-IP LTM for WebLogic

While we recommend using the application template, if you prefer to manually configure the BIG-IP LTM system, perform the following procedures.

To configure the BIG-IP LTM system to direct traffic to WebLogic Servers, you need to complete the following tasks:

- *Creating the HTTP health monitor*
- *Creating the pool*
- *Creating profiles*
- *Creating the virtual server*
- *Manually configuring the BIG-IP LTM to offload SSL (optional)*

Creating the HTTP health monitor

The first step is to set up health monitors for the WebLogic devices. This procedure is optional, but very strongly recommended. In our example, we create a simple HTTP health monitor. Although the monitor in the following example is quite simple, you can configure optional settings such as Send and Receive Strings to make the monitor much more specific.

To configure a health monitor

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**.
2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a name for the Monitor.
In our example, we type **bea-http**.
4. From the **Type** list, select **http**.
5. In the Configuration section, in the **Interval** and **Timeout** boxes, type an Interval and Timeout. We recommend at least a 1:3 +1 ratio between the interval and the timeout (for example, the default setting has an interval of **5** and an timeout of **16**). In our example, we use a **Interval** of **30** and a **Timeout** of **91**.
6. In the Send String and Receive Rule sections, you can add a Send String and Receive Rule specific to the device being checked.

The screenshot shows the 'New Monitor...' configuration window. The breadcrumb trail is 'Local Traffic >> Monitors >> New Monitor...'. The 'General Properties' section contains:

- Name: bea-http
- Type: HTTP
- Import Settings: http

 The 'Configuration' section is set to 'Basic' and includes:

- Interval: 30 seconds
- Timeout: 91 seconds
- Send String: GET /
- Receive String: (empty)
- User Name: (empty)
- Password: (empty)
- Reverse: Yes No
- Transparent: Yes No

 At the bottom are buttons for 'Cancel', 'Repeat', and 'Finished'.

Figure 2.1 *Creating the HTTP Monitor*

7. Click the **Finished** button.
The new monitor is added to the Monitor list.

Creating the pool

The first step is to define a load balancing pool for the WebLogic servers. A BIG-IP pool is a set of devices grouped together to receive traffic according to a load balancing method. This pool uses the monitor you just created.

To create a pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
The Pool screen opens.
2. In the upper right portion of the screen, click the **Create** button.
The New Pool screen opens.
3. In the **Name** box, type a name for your pool.
In our example, we use **bea-http-pool**.
4. In the **Health Monitors** section, select the name of the monitor you created in the *Creating the HTTP health monitor* section, and click the Add (<<) button. In our example, we select **bea-http**.

5. From the **Load Balancing Method** list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network).
In our example, we select **Least Connections (member)**.
6. In this pool, we leave the Priority Group Activation **Disabled**.
7. In the New Members section, make sure the **New Address** option button is selected.
8. In the **Address** box, add the first WebLogic server to the pool. In our example, we type **10.132.81.10**
9. In the **Service Port** box, type **7001**, the default service for WebLogic.
10. Click the **Add** button to add the member to the list.
11. Repeat steps 8-10 for each server you want to add to the pool.
In our example, we repeat these steps three times for the remaining servers, **10.132.81.11 - .13**.
12. Click the **Finished** button (see Figure 2.2).

Figure 2.2 Creating a pool for the WebLogic servers

Creating profiles

A *profile* is an object that contains user-configurable settings for controlling the behavior of a particular type of network traffic, such as HTTP connections. Using profiles enhances your control over managing network traffic, and makes traffic-management tasks easier and more efficient.

Although it is possible to use the default profiles, we strongly recommend you create new profiles based on the default parent profiles, even if you do not change any of the settings initially. Creating new profiles allows you to easily modify the profile settings specific to this deployment, and ensures you do not accidentally overwrite the default profile.

For the BEA WebLogic configuration, we create five new profiles: an HTTP profile, two TCP profiles, a persistence profile, and a OneConnect profile. If you plan on using the BIG-IP LTM system to offload SSL from the WebLogic devices, make sure to see *Creating a Client SSL profile*.

Creating an HTTP profile

In this procedure, we create an HTTP profile. The HTTP profile contains numerous configuration options for how the BIG-IP LTM system handles HTTP traffic. In the following example, we base our HTTP profile off of the **http-acceleration** parent profile, as we are using the WebAccelerator. If you are not using the WebAccelerator, we recommend using the **http-wan-optimized-compression-caching** parent.

To create a new HTTP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New HTTP Profile screen opens.
3. In the **Name** box, type a name for this profile. In our example, we type **bea-http-opt**.
4. From the **Parent Profile** list, select **http-acceleration** if you are using the WebAccelerator. If you are not using the WebAccelerator, select **http-wan-optimized-compression-caching**.
5. *Optional:* If you are using the BIG-IP LTM to offload SSL, in the Settings section, check the Custom box for **Redirect Rewrite**, and from the **Redirect Rewrite** list, select **Match**. See *Manually configuring the BIG-IP LTM to offload SSL*, on page 2-10 for more information.
6. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating the TCP profiles

The next profiles we create are the TCP profiles. If most of the BEA WebLogic users are accessing the devices via a Local Area Network, we recommend using the **tcp-lan-optimized** (for server-side TCP connections) parent profile. If the majority of the users are accessing the system from remote or home offices, we recommend using an additional TCP profile, called **tcp-wan-optimized** (for client side TCP connections). In our example, we leave these profiles at their default levels; you can configure any of the options as applicable for your network.

Creating the WAN optimized TCP profile

Now we configure the WAN optimized profile. Remember, if most of the users are accessing the system over the LAN or other low latency links, you do not need to create this profile.

To create a new TCP WAN optimized profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Protocol** menu, click **tcp**.
3. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **bea-tcp-wan**.
5. From the **Parent Profile** list, select **tcp-wan-optimized**.
6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.

Click the **Finished** button.

Creating the LAN optimized TCP profile

First we configure the LAN optimized profile. If you are not using version 9.4 or do not want to use this optimized profile, you can choose the default TCP parent profile.

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Protocol** menu, click **tcp**.
3. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **bea-tcp-lan**.
5. From the **Parent Profile** list, select **tcp-lan-optimized**.

6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating persistence profile

The next profile we create is a Persistence profile. We recommend using persistence for BEA WebLogic devices, although the type of persistence depends on your configuration. In our example, use cookie persistence (HTTP cookie insert).

To create a new cookie persistence profile based on the default profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, click **Persistence**. The Persistence Profiles screen opens.
3. In the upper right portion of the screen, click the **Create** button. The New Persistence Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **bea-cookie**.
5. From the **Persistence Type** list, select **Cookie**. The configuration options for cookie persistence appear.
6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button (see Figure 2.3).

General Properties	
Name	bea-cookie
Persistence Type	Cookie
Parent Profile	cookie

Configuration		Custom <input type="checkbox"/>
Cookie Method	HTTP Cookie Insert	<input type="checkbox"/>
Cookie Name		<input type="checkbox"/>
Expiration	<input checked="" type="checkbox"/> Session Cookie	<input type="checkbox"/>
Override Connection Limit	<input type="checkbox"/>	<input type="checkbox"/>

Cancel Repeat Finished

Figure 2.3 Creating the cookie persistence profile

Creating a OneConnect profile

The final profile we create is a OneConnect profile. With OneConnect enabled, client requests can utilize existing, server-side connections, thus reducing the number of server-side connections that a server must open to service those requests. This can provide significant performance improvements for WebLogic implementations. For more information on OneConnect, see the BIG-IP LTM documentation.

In our example, we leave all the options at their default settings. You can configure these options as appropriate for your network.

To create a new OneConnect profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Other** menu, click **OneConnect**. The Persistence Profiles screen opens.
3. In the upper right portion of the screen, click the **Create** button. The New HTTP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **bea-oneconnect**.
5. From the **Parent Profile** list, ensure that **oneconnect** is selected.
6. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating the virtual server

Next, we configure a virtual server that references the profiles and pool you created in the preceding procedures. If you are using the BIG-IP LTM system to offload SSL, be sure to see *Manually configuring the BIG-IP LTM to offload SSL*, on page 10 after completing this section.

To create the virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. In the upper right portion of the screen, click the **Create** button.
3. In the **Name** box, type a name for this virtual server. In our example, we type **bea-weblogic-http**.
4. In the **Destination** section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we use **192.168.10.100**.
6. In the **Service Port** box, type **80**.

General Properties	
Name	bea-weblogic-http
Destination	Type: <input checked="" type="radio"/> Host <input type="radio"/> Network Address: 192.168.10.100
Service Port	80 HTTP
State	Enabled

Figure 2.4 Creating the WebLogic virtual server

7. From the Configuration list, select **Advanced**.
The Advanced configuration options appear.
8. Leave the **Type** list at the default setting: **Standard**.
9. From the **Protocol Profile (Client)** list select the name of the profile you created in the *Creating the WAN optimized TCP profile* section. In our example, we select **bea-tcp-wan**.
10. From the **Protocol Profile (Server)** list, select the name of the profile you created in the *Creating the LAN optimized TCP profile* section. In our example, we select **bea-tcp-lan**.
11. From the **OneConnect Profile** list, select the name of the profile you created in *Creating a OneConnect profile*. In our example, we select **bea-oneconnect**.
12. From the HTTP Profile list, select the name of the profile you created in the *Creating an HTTP profile* section. In our example, we select **bea-http-opt**.

Configuration: Advanced	
Type	Standard
Protocol	TCP
Protocol Profile (Client)	bea-tcp-wan
Protocol Profile (Server)	bea-tcp-lan
OneConnect Profile	bea-oneconnect
HTTP Profile	bea-http-opt
FTP Profile	None

Figure 2.5 Selecting the BEA WebLogic profiles for the virtual server

13. In the Resources section, from the **Default Pool** list, select the pool you created in the *Creating the pool* section. In our example, we select **bea-http-pool**.

-
- From the **Default Persistence Profile** list, select the persistence profile you created in the *Creating persistence profile* section. In our example, we select **bea-cookie**.

Resources					
iRules	<table border="1"><thead><tr><th>Enabled</th><th>Available</th></tr></thead><tbody><tr><td></td><td><ul style="list-style-type: none">_sys_auth_ldap_sys_auth_radius_sys_auth_ssl_crdp</td></tr></tbody></table> <p>Up Down</p>	Enabled	Available		<ul style="list-style-type: none">_sys_auth_ldap_sys_auth_radius_sys_auth_ssl_crdp
Enabled	Available				
	<ul style="list-style-type: none">_sys_auth_ldap_sys_auth_radius_sys_auth_ssl_crdp				
HTTP Class Profiles	<table border="1"><thead><tr><th>Enabled</th><th>Available</th></tr></thead><tbody><tr><td></td><td><ul style="list-style-type: none">httpclass</td></tr></tbody></table> <p>Up Down</p>	Enabled	Available		<ul style="list-style-type: none">httpclass
Enabled	Available				
	<ul style="list-style-type: none">httpclass				
Default Pool	+ bea-http-pool				
Default Persistence Profile	bea-cookie				
Fallback Persistence Profile	None				

Cancel Repeat Finished

Figure 2.6 Adding the Pool and Persistence profile to the virtual server

- Click the **Finished** button.
The BIG-IP LTM HTTP configuration for the BEA WebLogic configuration is now complete.

Manually configuring the BIG-IP LTM to offload SSL

If you are using the BIG-IP LTM system to offload SSL from the BEA WebLogic devices, there are additional configuration procedures you must perform on the BIG-IP LTM system. In the following configuration, the BIG-IP LTM redirects all incoming traffic to the HTTP virtual server to the HTTPS virtual server. This is useful if a user types a URL in a browser, but forgets to change the protocol to HTTPS.

If your deployment does not require *all* traffic to be redirected to HTTPS, you do not need to configure the iRule or modify the HTTP virtual server as described below, nor configure the Rewrite Redirect setting in the HTTP profile in Step 6 of *Creating a new HTTP profile*, on page 11. You can have both an HTTP and HTTPS virtual server on the same address with the appropriate ports.

Important

This section is optional, and only necessary if you are using the BIG-IP LTM system for offloading SSL.

Using SSL certificates and keys

Before you can enable the BIG-IP LTM system to act as an SSL proxy, you must install a SSL certificate on the virtual server that you wish to use for BEA WebLogic connections on the BIG-IP LTM device. For this Deployment Guide, we assume that you already have obtained an SSL certificate, but it is not yet installed on the BIG-IP LTM system. For information on generating certificates, or using the BIG-IP LTM to generate a request for a new certificate and key from a certificate authority, see the **Managing SSL Traffic** chapter in the *Configuration Guide for Local Traffic Management*.

Importing keys and certificates

Once you have obtained a certificate, you can import this certificate into the BIG-IP LTM system using the Configuration utility. By importing a certificate or archive into the Configuration utility, you ease the task of managing that certificate or archive. You can use the Import SSL Certificates and Keys screen only when the certificate you are importing is in Privacy Enhanced Mail (PEM) format.

To import a key or certificate

1. On the Main tab, expand **Local Traffic**.
2. Click **SSL Certificates**. The list of existing certificates displays.
3. In the upper right corner of the screen, click **Import**.
4. From the **Import Type** list, select the type of import (Certificate or Key).

-
5. In the **Certificate** (or **Key**) **Name** box, type a unique name for the certificate or key.
 6. In the **Certificate** (or **Key**) **Source** box, choose to either upload the file or paste the text.
 7. Click **Import**.

If you imported the certificate, repeat this procedure for the key.

Creating a Client SSL profile

The next step in this configuration is to create a Client SSL profile. This profile contains the SSL certificate and Key information for offloading the SSL traffic.

To create a new Client SSL profile

1. On the Main tab, expand **Local Traffic**.
2. Click **Profiles**.
The HTTP Profiles screen opens.
3. On the Menu bar, from the SSL menu, select **Client**.
The Client SSL Profiles screen opens.
4. In the upper right portion of the screen, click the **Create** button.
The New Client SSL Profile screen opens.
5. In the **Name** box, type a name for this profile. In our example, we type **bea-clientssl**.
6. In the Configuration section, check the **Certificate** and **Key Custom** boxes.
7. From the **Certificate** list, select the name of the Certificate you imported in the *Importing keys and certificates* section.
8. From the **Key** list, select the key you imported in the *Importing keys and certificates* section.
9. Click the **Finished** button.

Creating a new HTTP profile

The next profile is a new HTTP profile that contains the necessary client header, along with the rewrite/redirect setting.

If you have already created an HTTP profile as described earlier in this guide, you can modify that profile with the modifications found in the following procedure.

To create a new HTTP profile for SSL

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New HTTP Profile screen opens.
3. In the **Name** box, type a name for this profile. In our example, we type **bea-ssl**.
4. From the **Parent Profile** list, select **http-acceleration** if you are using the WebAccelerator. If you are not using the WebAccelerator, select **http-wan-optimized-compression-caching**.
5. In the **Request Header Insert** row, check the Custom box. In the box, type: **WL-Proxy-SSL: true**.
6. In the **Redirect Rewrite** row, check the Custom box. From the **Redirect Rewrite** list, select **Match**.
7. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
8. Click the **Finished** button.

General Properties	
Name	bea-ssl
Parent Profile	http-acceleration

Settings		Custom <input type="checkbox"/>
Fallback Host		<input type="checkbox"/>
Fallback on Error Codes		<input type="checkbox"/>
Request Header Insert	WL-Proxy-SSL: true	<input checked="" type="checkbox"/>
Request Header Erase		<input type="checkbox"/>
Response Headers Allowed		<input type="checkbox"/>
Response Chunking	Selective	<input type="checkbox"/>
OneConnect Transformations	Enabled	<input type="checkbox"/>
Redirect Rewrite	Matching	<input checked="" type="checkbox"/>
Request Cookies		<input type="checkbox"/>

Figure 2.7 Creating the HTTP profile for SSL deployments

For more information on creating or modifying profiles, or applying profiles in general, see the BIG-IP documentation.

Creating the Redirect iRule

The Redirect iRule takes incoming HTTP requests (non-secure) and redirects them to the correct HTTPS (secure) virtual server, without user interaction.

To create the Redirect iRule

1. On the Main tab, expand **Local Traffic**, and then click **iRules**. The iRule screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New iRule screen opens.
3. In the **Name** box, enter a name for your iRule. In our example, we use **bea-httpstohttps**.
4. In the Definition section, copy and paste the following iRule:

```
when HTTP_REQUEST {  
    HTTP::redirect https://[HTTP::host][HTTP::uri]  
}
```
5. Click the **Finished** button.

Modifying the HTTP virtual server

The next task is to modify the HTTP virtual server you created in *Creating the virtual server*, on page 7 to use the iRule you just created.

To modify the existing BEA virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. From the Virtual Server list, click the WebLogic virtual server you created in the *Creating the virtual server* section. In our example, we click **bea-weblogic-vs**.
3. On the menu bar, click **Resources**. The Resources page for the virtual server opens.
4. From the **Default Pool** list, select **None**. This virtual server no longer requires the load balancing pool, as traffic is redirected to the HTTPS virtual server we create in the following procedure.
5. Click the **Update** button.
6. In the iRules section, click the **Manage** button. The Resource Management screen opens.
7. From the **Available** list, select the iRule you created in the *Creating the Redirect iRule* section, and click the Add (<<) button. In our example, we select **bea-httpstohttps**.
8. Click the **Finished** button.

Creating the HTTPS virtual server

The final task in this section is to create a HTTPS virtual server.

To create a new HTTPS virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New Virtual Server screen opens.
3. In the **Name** box, type a name for this virtual server. In our example, we type **bea-https-vs**.
4. In the **Destination** section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we use **192.168.104.147**.
6. In the **Service Port** box, type **443** or select **HTTPS** from the list.
7. From the Configuration list, select **Advanced**. The Advanced configuration options appear.
8. Leave the **Type** list at the default setting: **Standard**.
9. From the **Protocol Profile (Client)** list select the name of the profile you created in the *Creating the WAN optimized TCP profile* section. In our example, we select **bea-tcp-wan**.
10. From the **Protocol Profile (Server)** list, select the name of the profile you created in the *Creating the LAN optimized TCP profile* section. In our example, we select **bea-tcp-lan**.
11. From the **OneConnect Profile** list, select the name of the profile you created in *Creating a OneConnect profile*. In our example, we select **bea-oneconnect**.
12. From the HTTP Profile list, select the name of the profile you created in the *Creating a new HTTP profile* section. In our example, we select **bea-ssl**.
Make sure you have the Rewrite Redirect box checked in the HTTP profile as described in Step 5 of *Creating an HTTP profile*.
13. From the **SSL Profile (Client)** list, select the name of the SSL profile you created in the *Creating a Client SSL profile* section. In our example, we select **bea-clientssl**.
14. From the **Default Pool** list, select the pool you created in the *Creating the pool* section. In our example, we select **bea-http-pool**.
15. From the **Default Persistence Profile** list, select the persistence profile you created in the *Creating persistence profile*. In our example, we select **bea-cookie**.
16. Click the **Finished** button.



3

Manually Configuring the WebAccelerator with BEA WebLogic

- Creating an HTTP Class profile
- Modifying the Virtual Server to use the Class profile
- Creating an Application

Manually configuring the F5 WebAccelerator module for WebLogic

In this section, we configure the WebAccelerator module for the WebLogic devices to increase performance for end users. The F5 WebAccelerator is an advanced web application delivery solution that provides a series of intelligent technologies designed to overcome problems with browsers, web application platforms and WAN latency issues which impact user performance.

For more information on the F5 WebAccelerator, see www.f5.com/products/big-ip/product-modules/webaccelerator.html.

Prerequisites and configuration notes

The following are prerequisites for this section:

- ◆ We assume that you have already configured the BIG-IP LTM system for directing traffic to the BEA deployment as described in this Deployment Guide.
- ◆ You must have purchased and licensed the WebAccelerator module on the BIG-IP LTM system.
- ◆ You must have created an HTTP profile on the BIG-IP LTM system that has RAM Cache enabled. In our example (*Creating the HTTP profile*, on page 2-3) we use a parent profile that includes RAM Cache. If you did not create an HTTP profile with RAM Cache enabled, you must create a new HTTP profile, based on a parent profile that uses RAM Cache (we recommend HTTP Acceleration) and associate it with the virtual server.
- ◆ This document is written with the assumption that you are familiar with the BIG-IP LTM system, WebAccelerator and BEA WebLogic Server. Consult the appropriate documentation for detailed information.

Configuration example

Using the configuration in this section, the BIG-IP LTM system with WebAccelerator module is optimally configured to accelerate traffic to BEA WebLogic servers. The BIG-IP LTM with WebAccelerator module both increases end user performance as well as offloads the servers from serving repetitive and duplicate content.

In this configuration, a remote client with WAN latency accesses a BEA WebLogic server via the WebAccelerator. The user's request is accelerated on repeat visits by the WebAccelerator instructing the browser to use the dynamic or static object that is stored in its local cache. Additionally, dynamic and static objects are cached at the WebAccelerator so that they can be served quickly without requiring the server to re-serve the same objects.

Configuring the WebAccelerator module

Configuring the WebAccelerator module requires creating an HTTP class profile and creating an Application. The WebAccelerator device has a large number of other features and options for fine tuning performance gains, see the *WebAccelerator Administrator Guide* for more information.

Creating an HTTP Class profile

The first procedure is to create an HTTP class profile. When incoming HTTP traffic matches the criteria you specify in the WebAccelerator class, the system diverts the traffic through this class. In the following example, we create a new HTTP class profile, based on the default profile.

To create a new HTTP class profile

1. On the Main tab, expand **WebAccelerator**, and then click **Classes**. The HTTP Class Profiles screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New HTTP Class Profile screen opens.
3. In the **Name** box, type a name for this Class. In our example, we type **bea-class**.
4. From the Parent Profile list, make sure **httpclass** is selected.
5. In the Configuration section, from the **WebAccelerator** row, make sure **Enabled** is selected.
6. In the Hosts row, from the list select **Match Only**. The Host List options appear.
 - a) In the **Host** box, type the host name that your end users use to access the BEA WebLogic devices. In our example, we type **beaapplication.siterequest.com** (see Figure 3.1).
 - b) Leave the Entry Type at **Pattern String**.
 - c) Click the **Add** button.
 - d) Repeat these sub-steps for any other host names users might use to access the BEA deployment.
7. The rest of the settings are optional, configure them as applicable for your deployment.
8. Click the **Finished** button. The new HTTP class is added to the list.

Local Traffic >> Profiles : Protocol : HTTP Class >> New HTTP Class Profile...

General Properties

Name: bea-class

Parent Profile: httpclass

Configuration Custom

WebAccelerator: Enabled

Application Security: Disabled

Hosts: Match only...

Host List:

Host: beaapplication.siterequest.com

Entry Type: Pattern String

Add

beaapplication.siterequest.com

Delete

URI Paths: Match all

Headers: Match all

Cookies: Match all

Actions Custom

Send To: None

Rewrite URI:

Cancel Repeat Finished

Figure 3.1 Creating a new HTTP Class profile

Modifying the Virtual Server to use the Class profile

The next step is to modify the virtual server for your WebLogic deployment on the BIG-IP LTM system to use the HTTP Class profile you just created.

To modify the Virtual Server to use the Class profile

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. From the **Virtual Server** list, click the name of the virtual server you created for the WebLogic servers. In our example, we click **bea-weblogic-http**. The General Properties screen for the Virtual Server opens.
3. On the Menu bar, click **Resources**. The Resources screen for the Virtual Server opens.
4. In the HTTP Class Profiles section, click the **Manage** button.

5. From the **Available** list, select the name of the HTTP Class Profile you created in the preceding procedure, and click the Add (<<) button to move it to the Enabled box. In our example, we select **bea-class** (see Figure 3.2).
6. Click the **Finished** button. The HTTP Class Profile is now associated with the Virtual Server.

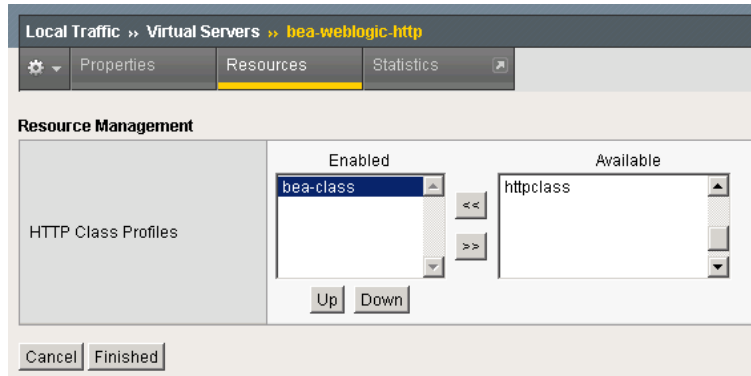


Figure 3.2 Adding the HTTP Class to the Virtual Server

◆ Important

*You must have created an HTTP profile on the BIG-IP LTM system that has RAM Cache enabled. In our example (**Creating an HTTP profile**, on page 2-4) we use a parent profile that includes RAM Cache. If you did not create an HTTP profile with RAM Cache enabled, you must create a new HTTP profile, based on a parent profile that uses RAM Cache (such as HTTP Acceleration), and modify the virtual server to use this new profile. This is only required for BIG-IP LTM version 9.4.2 and later.*

*To create the HTTP profile, use **Creating an HTTP profile**, on page 2-4, selecting the HTTP Acceleration parent profile. You must leave RAM Cache enabled; all other settings are optional. To modify the virtual server, follow Steps 1 and 2 from the preceding procedure to access the virtual server, and then from the HTTP Profile list, select the name of the new profile you just created and click Update.*

Creating an Application

The next procedure is to create a WebAccelerator Application. The Application provides key information to the WebAccelerator so that it can handle requests to your application appropriately.

To create a new Application

1. On the Main tab, expand **WebAccelerator**, and then click **Applications**.
The Application screen of the WebAccelerator UI opens in a new window.
2. Click the **New Application** button.
3. In the Application Name box, type a name for your application.
In our example, we type **BEA WebLogic**.
4. In the **Description** box, you can optionally type a description for this application.
5. From the **Local Policies** list, select **BEA WebLogic**. This is a pre-defined policy created specifically for BEA WebLogic devices (see Figure 4).
6. In the **Requested Host** box, type the host name that your end users use to access the BEA WebLogic deployment. This should be the same host name you used in Step 6a in the preceding procedure. In our example, we type **beaapplication.siterequest.com**.
If you have additional host names, click the **Add Host** button and enter the host name(s).
7. Click the **Save** button.

The screenshot shows the 'New Application' configuration page in the WebAccelerator UI. The breadcrumb navigation at the top reads 'Configuration > Applications > New Application'. The page is divided into several sections:

- General Options:** Contains an 'Application Name' field with the value 'BEA WebLogic' and a 'Description (optional)' text area containing the text 'WebAccelerator application for our Oracle/BEA WebLogic deployment'.
- Policies:** Contains two dropdown menus. 'Central Policy' is set to 'BEA Weblogic', and 'Remote Policy' is set to '- Select One -'.
- Hosts:** A table with two columns: 'Requested Host' and 'Action'. The 'Requested Host' column contains the value 'beaapplication.siterequest.com'. The 'Action' column contains links for 'Options' and 'Delete'.

At the bottom right of the form, there are three buttons: 'Add Host', 'Save' (highlighted in yellow), and 'Cancel'.

Figure 4 Configuring an Application on the WebAccelerator

The rest of the configuration options on the WebAccelerator are optional, configure these as applicable for your network. With this base configuration, your end users will notice a marked improvement in performance after their first visit.