

Deployment Guide

Deploying the BIG-IP LTM System with IBM WebSphere Servers



Deploying the BIG-IP LTM system and IBM WebSphere Servers

Welcome to the BIG-IP LTM system - IBM® WebSphere® Deployment Guide. This guide contains step-by-step procedures on how to configure the BIG-IP Local Traffic Management (LTM) system for directing traffic to the WebSphere servers. It also includes optional procedures for application optimization using BIG-IP compression and caching, as well using the BIG-IP LTM system for SSL termination.

IBM WebSphere Application Server provides the core software needed to deploy, integrate and manage e-business applications. F5 Networks BIG-IP LTM system is a secure, highly available and scalable application traffic management device.

This solution is powered in part by an iRule that enables persistence based on the application's own unique identifier (JSESSIONID).

For more information on IBM WebSphere, see <http://www-306.ibm.com/software/websphere/>

For more information on the BIG-IP LTM system, see <http://www.f5.com/products/big-ip/>.

Prerequisites and configuration notes

The following are prerequisites for this deployment:

- ◆ The BIG-IP LTM system must be running version 9.1 or later. We recommend version 9.2 or later.

For certain *optional* optimization features, the appropriate module must be licensed (such as compression and caching).

- ◆ This Deployment Guide was tested with IBM WebSphere 6.0.
- ◆ All of the configuration procedures in this document are performed on the BIG-IP LTM system. For information on how to deploy or configure IBM WebSphere, consult the appropriate IBM documentation.
- ◆ Briefly review the basic configuration tasks and the few pieces of information, such as IP addresses, that you should gather in preparation for completing this configuration.

◆ Note

This document is written with the assumption that you are familiar with both the BIG-IP LTM system and IBM WebSphere devices. For more information on configuring these products, consult the appropriate documentation.

Configuration example

Using the configuration in this guide, the BIG-IP LTM system is optimally configured to load balance traffic to IBM WebSphere servers. Figure 1 shows an example configuration with a redundant pair of BIG-IP devices and a cluster of WebSphere servers. In this configuration, we configure an iRule on the BIG-IP LTM system which uses the application's JSESSIONID for persistence.

◆ Tip

Although only one BIG-IP device is necessary for this configuration, we strongly recommend a redundant BIG-IP device for the highest level of availability.




Figure 1 BIG-IP WebSphere configuration example

◆ Note

The example in Figure 1 is a logical representation of this deployment. Your configuration may be dramatically different than the one shown.

Configuring the BIG-IP LTM system for deployment with IBM WebSphere Application Servers

To configure the BIG-IP LTM system for integration with IBM WebSphere Servers, you must complete the following procedures:

- *Connecting to the BIG-IP device*

-
- *Optional: Importing keys and certificates*
 - *Creating the HTTP health monitor*
 - *Creating the pool*
 - *Creating profiles*
 - *Creating the iRule*
 - *Creating the virtual servers*
 - *Synchronizing the BIG-IP configuration if using a redundant system*

◆ **Tip**

*We recommend you save your existing BIG-IP configuration before you begin the procedures in this Deployment Guide. To save your BIG-IP configuration, see **Appendix A: Backing up and restoring the BIG-IP LTM system configuration**, on page 17.*

The BIG-IP LTM system offers both web-based and command line configuration tools, so that users can work in the environment that they are most comfortable with. This Deployment Guide contains procedures to configure the BIG-IP LTM system using the BIG-IP web-based Configuration utility only. If you are familiar with using the **bigpipe** command line interface you can use the command line to configure the BIG-IP device, however, we recommend using the Configuration utility.

Connecting to the BIG-IP device

Use the following procedure to access the BIG-IP web-based Configuration utility using a web browser.

To connect to the BIG-IP LTM system using the Configuration utility

1. In a browser, type the following URL:
https://<administrative IP address of the BIG-IP device>
A Security Alert dialog box appears, click **Yes**.
The authorization dialog box appears.
2. Type your user name and password, and click **OK**.
The Welcome screen opens.

Once you are logged onto the BIG-IP LTM system, the Welcome screen of the new Configuration utility opens. From the Configuration utility, you can configure and monitor the BIG-IP LTM system, as well as access online help, download SNMP MIBs and Plug-ins, and even search for specific objects.

Optional: Importing keys and certificates

If you are using the BIG-IP LTM system for offloading SSL from the WebSphere devices, you must install a SSL certificate and key on the BIG-IP LTM system. 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 system 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*.

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

Important

If you are not using the BIG-IP LTM system for offloading SSL, you do not need to perform this procedure.

To import a key or certificate

1. On the Main tab, expand Local Traffic.
2. Click **SSL Certificates**.
This displays the list of existing certificates.
3. In the upper right corner of the screen, click **Import**.
4. From the **Import** list, select the type of import (**Key** or **Certificate**).
5. Select the import method (text or file).
6. Type the name of the key or certificate.
7. Click **Import**.

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

Creating the HTTP health monitor

The first step in this configuration is to set up an HTTP health monitor for the WebSphere devices. This procedure is optional, but very strongly recommended. For this configuration, we use an HTTP monitor, which checks nodes (IP address and port combinations), and can be configured to use **send** and **recv** statements in an attempt to retrieve explicit content from nodes. We configure the health monitors first in version 9.0 and later, as health monitors are now associated at the pool level.

To configure a health monitor from the BIG-IP Configuration utility

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**.
The Monitors screen opens.

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 **websphere_monitor**.
4. From the **Type** list, select **http**.
The HTTP Monitor configuration options appear.
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 an optional Send String and Receive Rule specific to the device being checked.
In our example, we are using the PlantsByWebSphere sample application, so we add a Send String of **GET /PlantsByWebSphere/ HTTP/1.0 \r\n\r\n**
And a Receive Rule of **<html>**. In this case, the monitor is successful if the opening HTML tag is returned.

General Properties	
Name	websphere_monitor
Type	HTTP
Import Settings	http
Configuration: Basic	
Interval	30 seconds
Timeout	91 seconds
Send String	GET /PlantsByWebSphere/ HTTP/1.0 \r\n\r\n
Receive String	<html>
User Name	
Password	
Reverse	<input type="radio"/> Yes <input checked="" type="radio"/> No
Transparent	<input type="radio"/> Yes <input checked="" type="radio"/> No
<input type="button" value="Cancel"/> <input type="button" value="Repeat"/> <input type="button" value="Finished"/>	

Figure 2 Creating the HTTP Monitor

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

◆ **Tip**

Although we strongly recommend a health monitor, it does not have to be an HTTP monitor. You can also configure multiple health monitors, such as configuring a basic TCP monitor in addition to the HTTP monitor.

Creating the pool

The next step in this configuration is to create a pool on the BIG-IP LTM system for the WebSphere devices. A BIG-IP pool is a set of devices grouped together to receive traffic according to a load balancing method.

To create the pool for the WebSphere devices

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.

***Note:** For more (optional) pool configuration settings, from the Configuration list, select **Advanced**. Configure these settings as applicable for your network.*

3. In the **Name** box, enter a name for your pool.
In our example, we use **websphere_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 **websphere_monitor**.
5. In the Resources section, 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 (node)**.
6. For 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 server to the pool. In our example, we type **10.133.22.14**
9. In the **Service Port** box, type the service number you want to use for this device, or specify a service by choosing a service name from the list.
In our example, we type **9080**, the default WebSphere port for the demonstration applications.
This may be different in your configuration.

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 twice for the remaining servers, **10.133.22.15** and **10.133.22.16**.
12. Click the **Finished** button (see Figure 3).

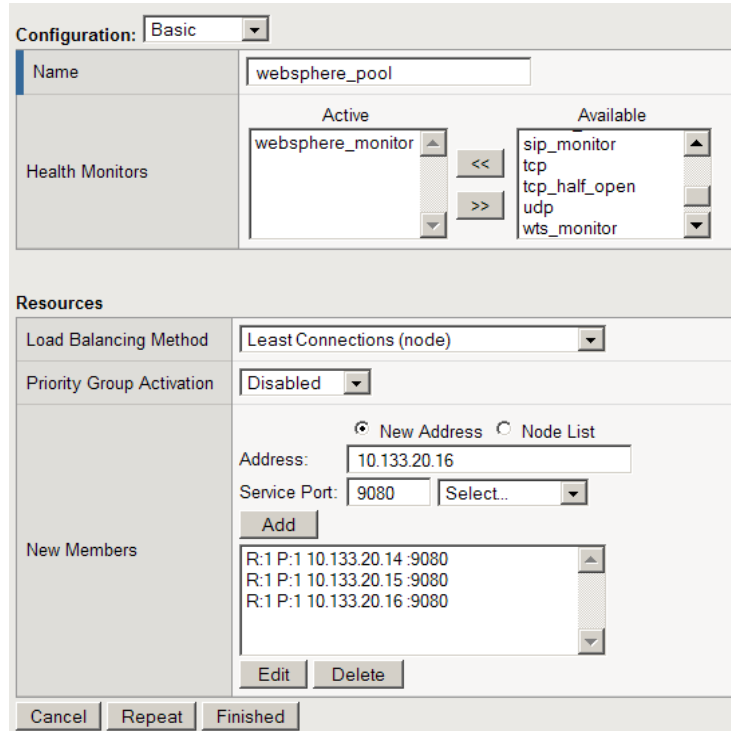


Figure 3 Creating the websphere_pool in the BIG-IP Configuration utility

Creating profiles

BIG-IP version 9.0 and later uses profiles. A *profile* is an object that contains user-configurable settings, with default values, 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. 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.

There are optional optimization settings you can configure if you want to optimize the BIG-IP LTM system for WebSphere deployments. These optional portions of the configuration will be clearly marked with *Optional Optimization*:

Creating an HTTP profile

The first new profile we create is an HTTP profile. If you are optimizing the BIG-IP configuration, the HTTP profile is where you configure the optional Intelligent Compression and Fast Cache options (these modules must be licensed on your BIG-IP LTM system). If you have not licensed a module, you will not see the options described in this procedure. Optimizing the HTTP profile provides the greatest improvement for WAN clients.

◆ Note

The following procedure shows one way to optimize the BIG-IP configuration, and was shown in our testing to give the greatest improvement. These procedures and the specific values given in some steps should be used as guidelines, modify them as applicable to your configuration.

To create an HTTP profile

1. On the Main tab, expand **Local Traffic**.
2. Click **Profiles**.
The HTTP 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 **websphere_http**.
5. Modify any of the settings as applicable for your network. If you are not using any of the Optional Optimization features, click the **Finished** button.

***Optional Optimization:** The following 8 steps are optional and show one way to optimally configure Compression on the BIG-IP LTM system. If your configuration does not contain compression, skip to Step 14.*

6. In the Settings table, from the **Response Chunking** section, click a check in the Custom box. From the list, select **Unchunk**. This allows for more efficient caching and compression.
7. In the Compression table, from the **Compression** row, click a check in the Custom box, then select **Enabled** from the list.
8. In the **Content list** section, we leave the settings at the default level, configure as applicable for your deployment.
9. In the **Compression Buffer Size** section, click a check in the Custom box. In the **Compression Buffer Size** box, type **131072**.

10. In the gzip **Compression Level** section, click a check in the Custom box. From the list, select a level of compression suitable to your configuration. For the most compression, select **9 - Most Compression (Slowest)**.
11. In the **gzip Memory Level** section, click a check in the Custom box. From the list, select **16** kilobytes.
12. In the **gzip Window size** section, click a check in the Custom box. From the list, select **64** kilobytes.
13. In the **HTTP/1.0 Requests** section, click a check in the Custom box. Click a check in the box to enable HTTP/1.0 requests.

Compression		Custom
Compression	Enabled	<input checked="" type="checkbox"/>
URI Compression	Not Configured	<input type="checkbox"/>
Content Compression	Content List...	<input type="checkbox"/>
Content List		
Content Type: <input type="text"/>		
<input type="button" value="Include"/> <input type="button" value="Exclude"/>		
Include List		
text/		
application/(xml x-javascript)		
application/(vnd.ms-excel vnd.ms-powerp		
Exclude List		
<input type="button" value="Edit"/> <input type="button" value="Delete"/>		
Preferred Method	Gzip	<input type="checkbox"/>
Minimum Content Length	1024 bytes	<input type="checkbox"/>
Compression Buffer Size	131072 bytes	<input checked="" type="checkbox"/>
gzip Compression Level	9 - Most Compression (Slowest)	<input checked="" type="checkbox"/>
gzip Memory Level	16 kilobytes	<input checked="" type="checkbox"/>
gzip Window Size	64 kilobytes	<input checked="" type="checkbox"/>
Vary Header	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/>
HTTP/1.0 Requests	<input checked="" type="checkbox"/> Enabled	<input checked="" type="checkbox"/>
Keep Accept Encoding	<input type="checkbox"/>	<input type="checkbox"/>
Browser Workarounds	<input type="checkbox"/>	<input type="checkbox"/>
CPU Saver	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/>
CPU Saver High Threshold	90 %	<input type="checkbox"/>
CPU Saver Low Threshold	75 %	<input type="checkbox"/>

Figure 4 Optional Compression configuration settings

Optional Optimization: The following 8 steps are optional and show one way to optimally configure Caching on the BIG-IP LTM system. If your configuration does not include caching on the BIG-IP LTM system, click the **Finished** button.

Note: Implementing Caching on the BIG-IP LTM system requires that cache control headers be issued by the WebSphere Application, or you must add an entry in the **URI List** to exclude parts of the application that should not be cached. In our example, we added an entry to the Exclude list to avoid caching the Shopping Cart. See Step 20.

14. In the **RAM Cache** table, click a check in the Custom boxes for all the settings *except* Maximum Entries, URI Caching, Ignore Headers, and Aging Rate.
15. In the **Ram Cache** section, select **Enabled** from the list.
16. In the **Maximum Cache Size** section, type **10** in the box.
17. In the **Maximum Age** section, type **86400** seconds.
18. In the **Minimum Object Size** section, type **0**.
19. In the **Maximum Object Size** section, type **2,000,000** bytes.
20. *Optional* (see the **Note** above): In the URI List section, in the **URI** box, type a URI to exclude, and click the **Exclude** button.
In our example, we type **/PlantsByWebSphere/servlet/ShoppingServlet** to avoid caching the Shopping Cart, and click the **Exclude** button.
21. In the **Insert Age Header** section, select Disabled from the list.
22. Click the **Finished** button.

The screenshot shows the 'RAM Cache' configuration window with a 'Custom' tab selected. The settings are as follows:

Setting	Value	Custom
RAM Cache	Enabled	<input checked="" type="checkbox"/>
Maximum Cache Size	10 megabytes	<input checked="" type="checkbox"/>
Maximum Entries	10000	<input type="checkbox"/>
Maximum Age	86400 seconds	<input checked="" type="checkbox"/>
Minimum Object Size	0 bytes	<input checked="" type="checkbox"/>
Maximum Object Size	2000000 bytes	<input checked="" type="checkbox"/>
URI Caching	URI List...	<input checked="" type="checkbox"/>
URI List	URI: <input type="text"/> Pin <input type="button" value="Pin"/> Include <input type="button" value="Include"/> Exclude <input type="button" value="Exclude"/> Pin List <input type="text"/> Include List <input type="text"/> Exclude List <input type="text" value="/PlantsByWebSphere/servlet/ShoppingServlet"/> <input type="button" value="Edit"/> <input type="button" value="Delete"/>	
Ignore Headers	All	<input type="checkbox"/>
Insert Age Header	Disabled	<input checked="" type="checkbox"/>
Aging Rate	9	<input type="checkbox"/>

Figure 5 Optional RAM Cache configuration settings

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

Creating a TCP profile

For this Deployment Guide, we create a new TCP profile based off of the default TCP profile. This new profile allows you to change the default setting for TCP connection timeout and other options.

The BIG-IP LTM system's TCP Express feature set provides a number of enhancements and optimizations to TCP handling that enhance end user experience. These settings are found in the Optional Optimizations steps.

To create a new TCP profile based on the default TCP profile

1. On the Main tab, expand **Local Traffic**.
2. Click **Profiles**.
The HTTP Profiles screen opens.
3. On the Menu bar, from the **Protocol** menu, select **TCP**.
4. In the upper right portion of the screen, click the **Create** button.
The New TCP Profile screen opens.
5. In the **Name** box, type a name for this profile. In our example, we type **websphere_tcp**.
6. Modify any of the settings as applicable for your network.

***Optional Optimization:** The following 4 steps are optional and show one way to optimally configure the TCP profile.*

7. In the Configuration table, locate **Proxy Buffer Low**, and click a check in the Custom box on the far right. In the Proxy Buffer Low box, type **131072**.
8. In the **Proxy Buffer High** section, click a check in the Custom box, and in the Proxy Buffer High box, type **131072**.
9. In the **Send Buffer** section, click a check in the Custom box, and in the Send Buffer box, type **65535**.
10. In the **Receive Window** section, click a check in the Custom box, and in the Receive Window box, type **65535**.
11. Click the **Finished** button (see Figure 6).

General Properties	
Name	websphere_tcp
Parent Profile	tcp
Settings Custom	
Reset On Timeout	<input checked="" type="checkbox"/> Enabled <input type="checkbox"/>
Time Wait Recycle	<input checked="" type="checkbox"/> Enabled <input type="checkbox"/>
Delayed Acks	<input checked="" type="checkbox"/> Enabled <input type="checkbox"/>
Proxy Maximum Segment	<input type="checkbox"/> <input type="checkbox"/>
Proxy Options	<input type="checkbox"/> <input type="checkbox"/>
Proxy Buffer Low	131072 bytes <input checked="" type="checkbox"/>
Proxy Buffer High	131072 bytes <input checked="" type="checkbox"/>
Idle Timeout	Specify... 1200 seconds <input checked="" type="checkbox"/>
Time Wait	Specify... 2000 seconds <input type="checkbox"/>
Fin Wait	Specify... 5 seconds <input type="checkbox"/>
Close Wait	Specify... 5 seconds <input type="checkbox"/>
Send Buffer	65535 bytes <input checked="" type="checkbox"/>
Receive Window	65535 bytes <input checked="" type="checkbox"/>
Keep Alive Interval	Specify... 1800 seconds <input type="checkbox"/>
Maximum Syn	4 <input type="checkbox"/>

Figure 6 Optional settings in the TCP profile

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

Creating a Client SSL profile

If you are using the BIG-IP LTM system to offload SSL, you must create an Client SSL profile. This profile contains the SSL certificate and Key information for offloading the SSL traffic.

◆ Note

If you are not using the BIG-IP LTM system for offloading SSL, you do not need to create this profile.

To create a new Client SSL 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, from the **SSL** menu, select **Client**.
The Client SSL Profiles screen opens.
 3. In the upper right portion of the screen, click the **Create** button.
The New Client SSL Profile screen opens.
 4. In the **Name** box, type a name for this profile. In our example, we type **websphere_clientSSL**.
 5. In the Configuration section, click a check in the **Certificate** and **Key Custom** boxes.
 6. From the **Certificate** list, select the name of the Certificate you imported in the *Optional: Importing keys and certificates* section.
 7. From the **Key** list, select the key you imported in the *Optional: Importing keys and certificates* section.
 8. Click the **Finished** button.

For more information on SSL certificates, or creating or modifying profiles, see the BIG-IP documentation.

Creating a OneConnect Profile

The final profile we create is a OneConnect™ profile. OneConnect improves performance by aggregating multiple client requests into a server-side connection pool, enabling client requests to reuse server-side connections. For more information on the OneConnect profile, see the *BIG-IP Configuration Guide for Local Traffic Management*.

To create a 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, select **OneConnect**.
The OneConnect profile screen opens.
3. In the upper right portion of the screen, click the **Create** button.
The New OneConnect Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **websphere_oneconnect**.
5. Modify any of the settings as applicable for your configuration. In our example, we leave the settings at their default levels.
6. Click the **Finished** button.

Creating the iRule

The next step is to configure an iRule on the BIG-IP LTM system that allows the BIG-IP LTM system to use the application's JSESSIONID for persistence. The iRule looks for the JSESSIONID in the cookie, but also checks the URI if the cookie does not exist.

◆ Note

If your configuration does not require persistence, this procedure is not necessary.

To create the 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 **WebSphereJsessionID**.
4. In the Definition section, copy and paste the following iRule.

```
when CLIENT_ACCEPTED {
  set add_persist 1
}
when HTTP_RESPONSE {
  if { [HTTP::cookie exists "JSESSIONID"] and $add_persist } {
    persist add uie [HTTP::cookie "JSESSIONID"]
    set add_persist 0
  }
}
when HTTP_REQUEST {
  if { [HTTP::cookie exists "JSESSIONID"] } {
    persist uie [HTTP::cookie "JSESSIONID"]
  } else {
    set jsess [findstr [HTTP::uri] "jsessionid" 11 ";"]
    if { $jsess != "" } {
      persist uie $jsess
    }
  }
}
```

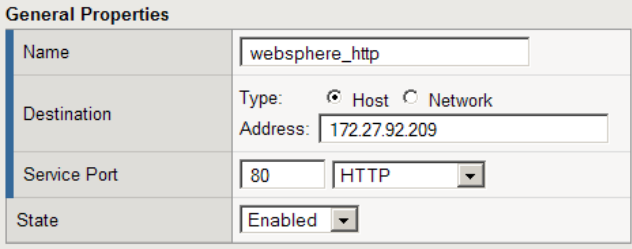
5. Click the **Finished** button.

Creating the virtual servers

Next, we configure the virtual servers. In our example, we create two virtual servers, one for port **80** and one for port **9080**. We created two virtual servers because in our testing, the WebSphere application redirected traffic to port **9080**. This redirection was observed when the full URI was not used, so we added the port 9080 virtual server to catch those instances.

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.
The New Virtual Server screen opens.
3. In the **Name** box, type a name for this virtual server. In our example, we type **websphere_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 **172.27.92.209**.
6. In the **Service Port** box, type **80** (see Figure 5).



General Properties	
Name	websphere_http
Destination	Type: <input checked="" type="radio"/> Host <input type="radio"/> Network
	Address: 172.27.92.209
Service Port	80 HTTP
State	Enabled

Figure 7 General Properties of the virtual server

7. From the Configuration list, select **Advanced**.
The Advanced options appear.
8. From the **Protocol Profile (Client)** list select the name of the profile you created in the *Creating a TCP profile* section. In our example, we select **websphere_tcp**.
9. 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 **websphere_http**.
10. If you are using the BIG-IP LTM system for offloading SSL, and created an SSL profile, from the **SSL Profile (Client)** list, select the name of the profile you created in the *Creating a Client SSL profile* section. In our example, we select **websphere_clientSSL** (see Figure 8).

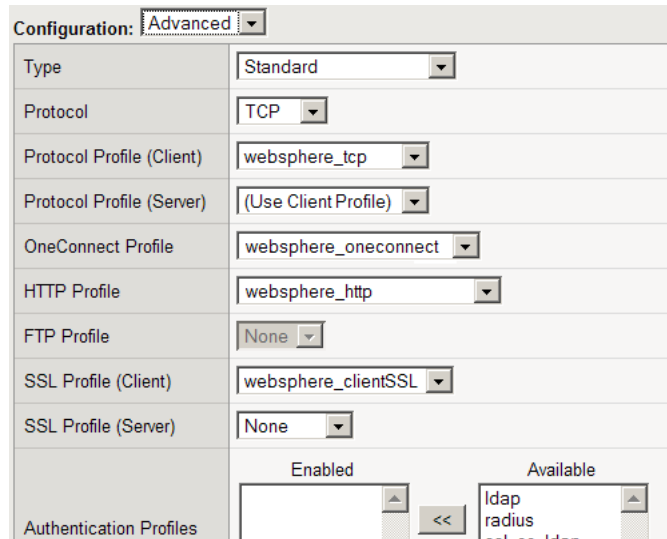


Figure 8 Selecting newly created objects for the virtual server

11. In the Resources section, from the **iRules** Available list, select the name of the iRule you created in the *Creating the iRule* section, and click the Add (<<) button to add it to the Enabled box. In our example, we enable **WebSphereJsessionID**.
12. From the **Default Pool** list, select the pool you created in the *Creating the pool* section. In our example, we select **webspere_pool**.
13. Click the **Finished** button.

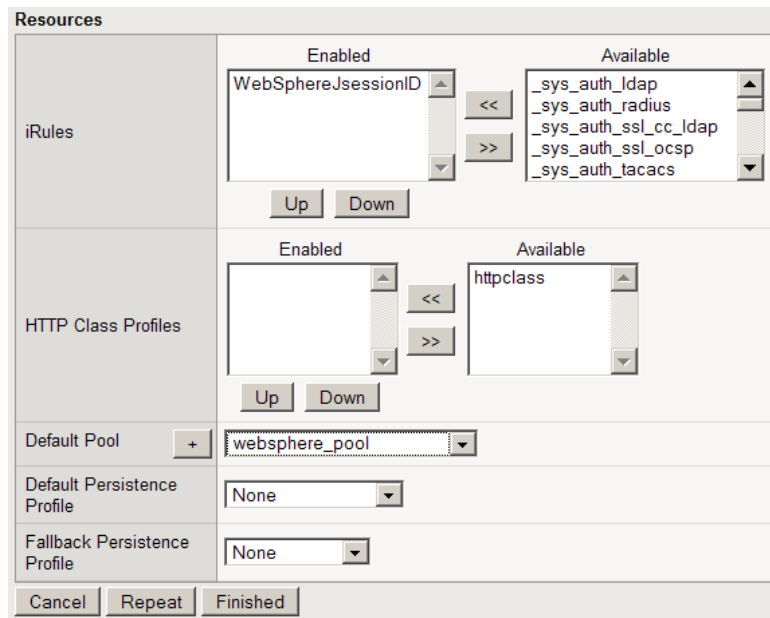


Figure 9 Resources section of the add virtual server page

To create the virtual server for port 9080

Repeat the preceding procedure for the second virtual server, with the following exceptions:

- In Step 3, type a different name for this virtual server. In our example, we use **websphere9080_virtual**.
- In Step 6, type **9080** for the port.

All of the rest of the steps are the same.

Synchronizing the BIG-IP configuration if using a redundant system

If you are using a redundant BIG-IP configuration, the final step is to synchronize the configuration to the peer BIG-IP device.

To synchronize the configuration using the Configuration utility

1. On the Main tab, expand **System**.
2. Click **High Availability**.
The Redundancy screen opens.
3. On the Menu bar, click **ConfigSync**.
4. Click the **Self --> Peer** button.
The configuration synchronizes with its peer.

Appendix A: Backing up and restoring the BIG-IP LTM system configuration

We recommend saving your BIG-IP configuration before you begin this configuration. When you save the BIG-IP configuration, it collects the following critical data and compress it into a single User Configuration Set (UCS) file:

- BIG-IP configuration files
- BIG-IP license and passwords
- SSL certificates
- SSH keys

Saving up and restoring the BIG-IP configuration

The Configuration Management screen allows you to save and restore all configuration files that you may edit to configure a BIG-IP LTM system. These configuration files are called a User Configuration Set (UCS). The Configuration Management screen contains sections for saving and restoring a configuration. The list boxes in these sections display only files in the `/usr/local/ucs` directory. If you want to save or restore files from another directory, you must type the full path in the box.

To save the BIG-IP configuration using the Configuration utility

1. In the navigation pane, click **System Admin**.
The User Administration screen displays.
2. Click the Configuration Management tab.
The Configuration Management screen displays.
3. In the **Save Current Configuration** section, type the path where you want your configuration file saved or choose a path from the list box. If no path is specified, the BIG-IP saves files to `/usr/local/ucs`. The BIG-IP appends the extension `.ucs` to file names without it. In our example, we type `pre_webisphere_backup.ucs`.
4. Click the **Save** button to save the configuration file.

To restore a BIG-IP configuration

1. In the navigation pane, click **System Admin**.
The User Administration screen displays.
2. Click the Configuration Management tab.
The Configuration Management screen displays.
3. In the **Restore a Configuration** section, choose the configuration file you want to restore from the list box, or type the path where your configuration files were saved.
4. Click the **Restore** button.
To check the status of the restoration, click the **View Log** button. You should wait a few moments for the log file to start generating before you click **View Log**. Repeated clicking of this button will update your screen with the most current log file information until the restoration is complete.