DEPLOYMENT GUIDE
Version 1.1

Deploying F5 with IBM WebSphere 7

Important: This guide has been archived. While the content in this guide is still valid for the products and versions listed in the document, it is no longer being updated and may refer to F5 or third party products or versions that have reached end-of-life or end-of-support. For a list of current guides, see https://f5.com/solutions/deployment-guides.

Table of Contents

1

Deploying the BIG-IP LTM with IBM WebSphere 7

- Configuring the BIG-IP LTM with IBM HTTP Server and the WebSphere plugin
- Creating the HTTP health monitor
- Creating the IBM HTTP server pool
- Creating profiles
- Creating the Redirect iRule
- Creating the virtual 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 Manager (LTM) for directing traffic to the IBM HTTP servers and WebSphere application servers.

IBM WebSphere provides software for SOA environments that enables dynamic, interconnected business processes, and delivers highly effective application infrastructures for all business situations.

For more information on IBM WebSphere, see **http://www-01.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 10.0.1 or later.
- For this guide, we assume that the WebSphere implementation is already deployed and configured properly. This deployment guide was written so you would not have to make changes to the application.
- In this guide, we use the Plants by WebSphere sample application.
- This Deployment Guide was tested with IBM WebSphere 7.0.5. All of the configuration procedures in this document are performed on F5 devices. 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.

Product versions and revision history

Product and versions tested for this deployment guide:

Product Tested	Version Tested
BIG-IP LTM and WebAccelerator	v10.0.1 (applicable to v9.4.7 and later)
FirePass Controller	v6.0.2
IBM WebSphere	7.0.5

Revision history:

Document Version	Description
1.0	New deployment guide
1.1	Added support for BIG-IP v10.1

Configuration example

Using the configuration in this guide, the BIG-IP LTM system is optimally configured to load balance traffic to IBM HTTP Servers with the WebSphere plugin. Figure 1.1 shows an example configuration with a redundant pair of BIG-IP devices and a cluster of WebSphere servers.

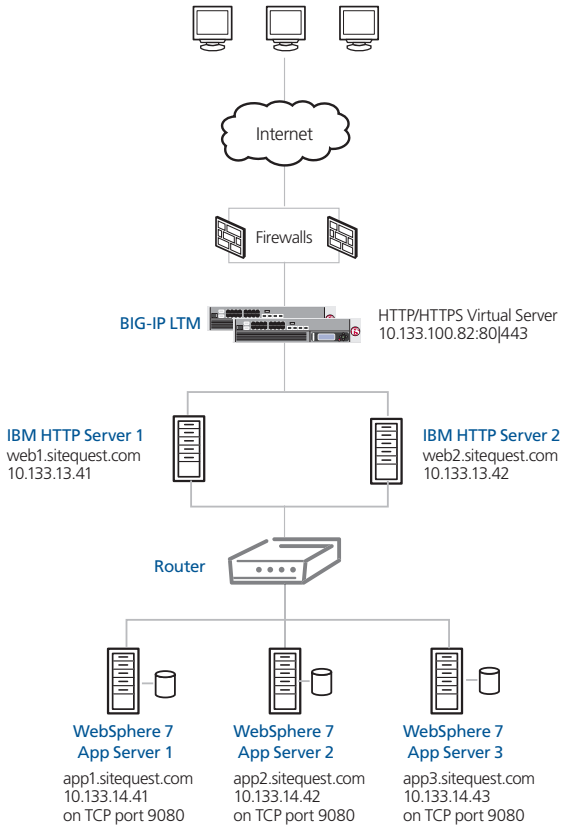

Figure 1.1 *BIG-IP WebSphere configuration example*

Configuring the BIG-IP LTM with IBM HTTP Server and the WebSphere plugin

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 IBM HTTP server pool*
- *Creating profiles*
- *Creating the Redirect iRule*
- *Creating the virtual servers*

Tip

We recommend you save your existing BIG-IP configuration before you begin the procedures in this Deployment Guide. For information on backing up or restoring a BIG-IP LTM configuration, refer to the appropriate BIG-IP LTM manual, available on ***Ask F5****.*

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 IBM 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**.
8. 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. This monitor is used for health checking both the web server and application server WebSphere components. 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 Receive strings in an attempt to retrieve explicit content from nodes, as we show in the following example.

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.

To configure the 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 **ibm-web-monitor**.
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** box, you can add an optional Send String specific to the device or application 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.
7. In the **Receive String** box, type what you expect the server to return as a result of the Send String. In our example, we type **<html>**: the monitor is successful if the opening HTML tag is returned.

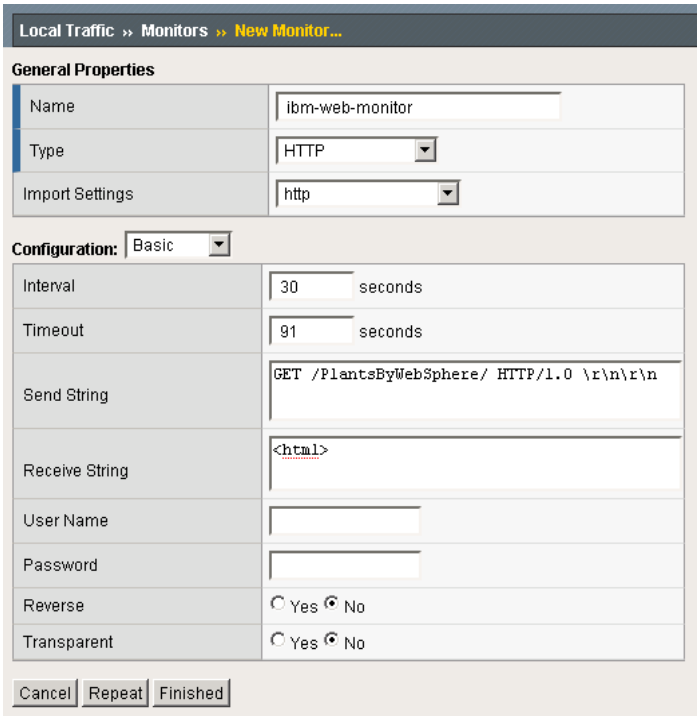

Figure 1.2 *Creating the HTTP Monitor*

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

Creating the IBM HTTP server pool

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

To create the pool for the HTTP servers

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. From the Configuration list, select **Advanced**.
4. In the **Name** box, enter a name for your pool.
   In our example, we use **ibm-web-pool**.
5. 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 **ibm-web-monitor**.
6. *Optional*: In the **Slow Ramp Time** box, type **300**.
   Because we are using the Least Connections load balancing method, we set the Slow Ramp Time in order to ensure that if a pool member becomes available after maintenance or a new member is added, the BIG-IP does not send all new connections to that member (a newly available member always has the least number of connections).

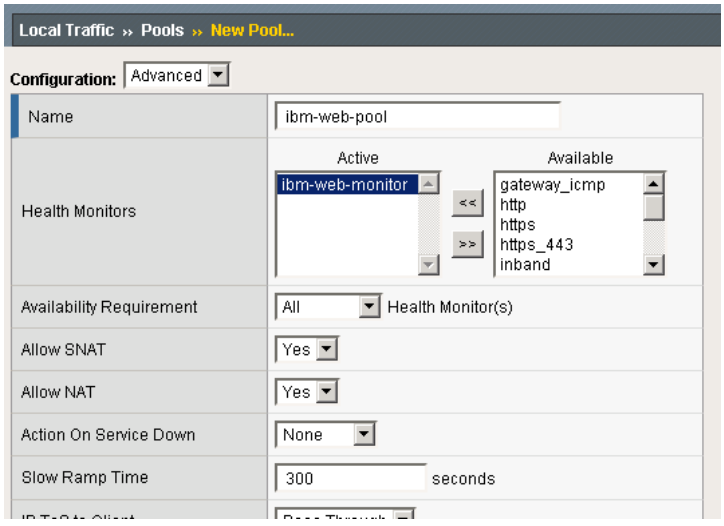

Figure 1.3 *Configuring the BIG-IP pool (truncated)*

7. 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 (member)**.

Note: *If you are running multiple applications on the same physical WebSphere application server, we recommend you use* ***Least Connections (Node)****.*

8. For this pool, we leave the Priority Group Activation **Disabled**.
9. In the New Members section, make sure the **New Address** option button is selected.
10. In the **Address** box, add the first server to the pool. In our example, we type **10.133.13.41**.
11. In the **Service Port** box, type **80** or select **HTTP** from the list.
12. Click the **Add** button to add the member to the list.
13. Repeat steps 9-11 for each server you want to add to the pool. In our example, we repeat these steps one time for the remaining server, **10.133.13.42**.
14. Click the **Finished** button.

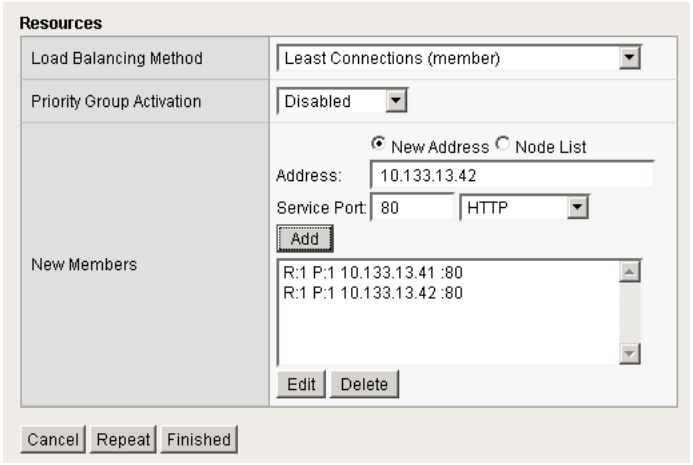

Figure 1.4 *Resources section of the IBM HTTP Server pool*

Creating profiles

The next task is to configure the 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.

Creating an HTTP profile

The first new profile we create is an HTTP profile. The HTTP profile contains numerous configuration options for how the BIG-IP LTM system handles HTTP traffic. For deployments where the majority of users accessing the WebSphere devices are connecting across a WAN, F5 recommends enabling compression and caching on the BIG-IP LTM by using the **http-wan-optimized-compression-caching**. This profile uses specific compression and caching (among other) settings to optimize traffic over the WAN.

Important

*If you are using the BIG-IP with the WebAccelerator module, use the **http-acceleration** parent profile.*

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 **ibm-http-opt**.
4. From the **Parent Profile** list, select one of the following:
   - If you are not using the WebAccelerator, select **http-wan-optimized-compression-caching**.
   - If you are using the BIG-IP WebAccelerator, select the **http-acceleration**.
5. *Optional:* If you 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**.
6. *Optional*: 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.
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.

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

Creating a TCP profile

The next profiles we create are the TCP profiles. If most of the IBM WebSphere 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.

For the internal forwarding virtual servers, we recommend creating an additional TCP LAN profile.

Creating the LAN optimized TCP profile

First we configure the LAN optimized profile. If you 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 **ibm-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 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 **ibm-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.
7. Click the **Finished** button.

Creating the persistence profiles

Next, we create the persistence profiles. We recommend using cookie persistence (HTTP cookie insert) as the default profile for the front end HTTP servers, and configuring Source Address persistence as a fallback mode.

Creating the cookie persistence profile

Use this procedure to configure the cookie persistence profile.

To create a new cookie persistence 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 **ibm-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.

Creating the source address persistence profile

Use this procedure to configure the source address persistence profile.

To create a new source address persistence profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, click **Persistence**.
3. Click the **Create** button.
4. In the **Name** box, type a name for this profile. In our example, we type **ibm-source**.
5. From the **Persistence Type** list, select **Source Address Affinity**.
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 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 **ibm-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 OneConnect, see SOL7208 (https://support.f5.com/kb/en-us/solutions/public/7000/200/sol7208.html) on Ask F5.

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 **ibm-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 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 **ibm-httptohttps**.
4. In the Definition section, type the following iRule (you can also copy and paste, but remove the line numbers):

```
1  when HTTP_REQUEST {
2  HTTP::redirect https://[HTTP::host][HTTP::uri]
3  }
```

5. Click the **Finished** button.

Creating the virtual servers

Next, we configure the virtual servers. A virtual server with its virtual address is the visible, routable entity through which the WebSphere devices in a load balancing pool are made available to the client (the IP address to give clients or add to DNS).

Creating the HTTP server virtual server

The first virtual server we create is the HTTP virtual server. If you are using the BIG-IP LTM to offload SSL, this virtual server is solely to intercept incoming HTTP traffic and redirect it to HTTPS using the iRule you

created; this virtual is optional. The second virtual server terminates the SSL (HTTPS) connections and sends traffic via HTTP to the pool of IBM devices.

To create the HTTP 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 **ibm-http-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 **10.133.100.132**.
6. In the **Service Port** box, type **80**, or select **HTTP** from the list.
7. In the Configuration section, select **Advanced** from the list. The Advanced configuration options appear.
8. 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 **ibm-tcp-wan**.
   This is optional, an only necessary if you created a WAN optimized profile.
9. 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 **ibm-tcp-lan**.
10. 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 **ibm-http-opt**.
11. If you are using the BIG-IP LTM to offload SSL:
    In the Resources section, from the **iRules** Available list, select the iRule you created for redirection in the *Creating the Redirect iRule* section. In our example, we select **ibm-httptohttps**.
12. If you are ***not*** using the BIG-IP LTM to offload SSL:
    a) From the **Default Pool** list, select the name of the pool you created in *Creating the IBM HTTP server pool*, on page 1-6.
    If you using the BIG-IP to offload SSL, do not select a pool.
    b) From the **Default Persistence Profile** list, select the persistence profile you created in the *Creating the cookie persistence profile* section. In our example, we select **ibm-cookie**.
    c) From the **Fallback Persistence Profile** list, select the persistence profile you created in the *Creating the source address persistence profile* section. In our example, we select **ibm-source**.

13. Click the **Finished** button.

Creating the HTTPS virtual server

Next, we create the HTTPS virtual server.

To create the HTTPS virtual server

1. On the Main tab, expand **Local Traffic**, click **Virtual Servers**, and then click the **Create** button. The New Virtual Server screen opens.
2. In the **Name** box, type a name for this virtual server. In our example, we type **ibm-https-vs**.
3. 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 **10.133.100.82**.
4. In the **Service Port** box, type **443**, or select **HTTPS** from the list.

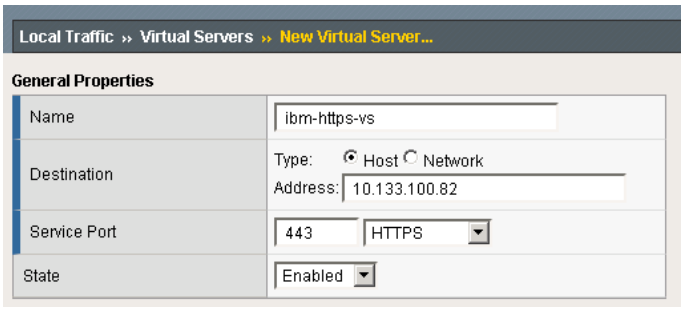

Figure 1.5 *Configuring the virtual server general properties*

5. From the Configuration list, select **Advanced**.
6. 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 **ibm-tcp-wan**. This is optional.
7. 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 **ibm-tcp-lan**.
8. From the **OneConnect Profile** list, select the name of the profile you created in the *Creating a OneConnect Profile* section. In our example, we select **ibm-oneconnect**.
9. From the **HTTP Profile** list, select the name of the profile you created in the *Creating a TCP profile* section. In our example, we select **ibm-http-opt**.
10. From the **SSL Profile (Client)** list, select the SSL profile you created in the *Creating a Client SSL profile* section. In our example, we select **ibm-clientssl**.

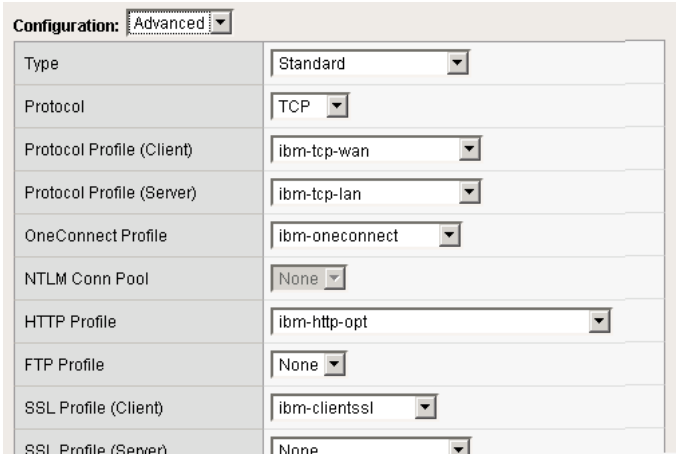

Figure 1.6 *Configuration section of the virtual server (condensed)*

11. In the Resources section, from the **Default Pool** list, select the pool you created in the *Creating the IBM HTTP server pool* section. In our example, we select **ibm-web-pool**.
12. From the **Default Persistence Profile** list, select the persistence profile you created in the *Creating the cookie persistence profile* section. In our example, we select **ibm-cookie**.
13. From the **Fallback Persistence Profile** list, select the persistence profile you created in the *Creating the source address persistence profile* section. In our example, we select **ibm-source**.
14. Click the **Finished** button.

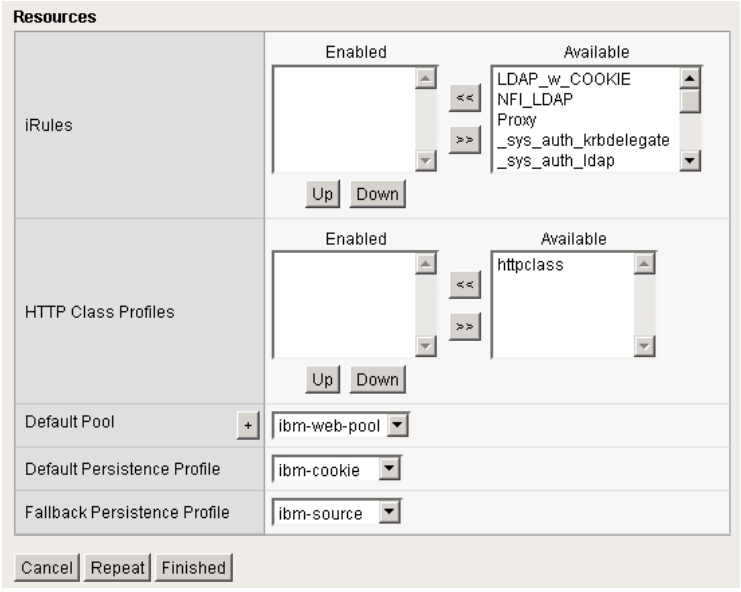

Figure 1.7 *Resources section of the virtual server*

This concludes the BIG-IP LTM configuration for the IBM HTTP Servers. For advanced load balancing and health monitoring configuration procedures, see Chapter 3, *Configuring advanced load balancing and health monitoring*.

2

Deploying the BIG-IP WebAccelerator with IBM WebSphere 7

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

Configuring the F5 WebAccelerator module with IBM WebSphere 7.0

In this section, we configure the WebAccelerator module for the WebSphere 7.0 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 WebSphere deployment as described in this Deployment Guide.
- You must have purchased and licensed the WebAccelerator module on the BIG-IP LTM system, version 9.4 or later.
- The BIG-IP LTM must have an HTTP profile that has RAM Cache enabled. If you did not create an HTTP profile with RAM Cache enabled, return to *Creating an HTTP profile*, on page 1-8, and create a new HTTP profile based on the **http-acceleration** parent profile.
- This document is written with the assumption that you are familiar with the BIG-IP LTM system, WebAccelerator and IBM WebSphere. 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 IBM WebSphere 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 WebSphere 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, for more information, see the ***WebAccelerator Administrator Guide*** available on ***Ask F5***.

Connecting to the BIG-IP LTM device

Use the following procedure to access the BIG-IP LTM system's 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.

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 **Class Profiles**.
   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 **websphere-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 WebSphere devices. In our example, we type **websphere.siterequest.com**.
   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 WebSphere 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.

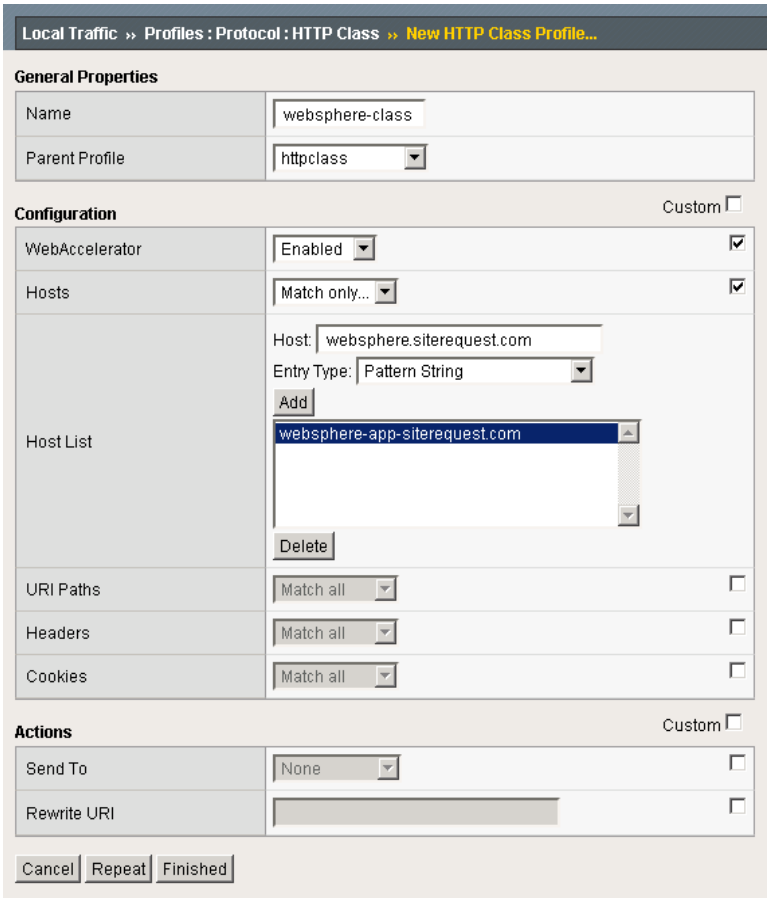

Figure 2.1 *Creating a new HTTP Class profile*

Modifying the Virtual Server to use the Class profile

The next step is to modify the BIG-IP LTM virtual server for your WebSphere deployment 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 IBM HTTP servers. In our example, we click **ibm-http-vs**. The General Properties screen opens.
3. On the Menu bar, click **Resources**.
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 **websphere-class**.
6. Click the **Finished** button. The HTTP Class Profile is now associated with the Virtual Server.

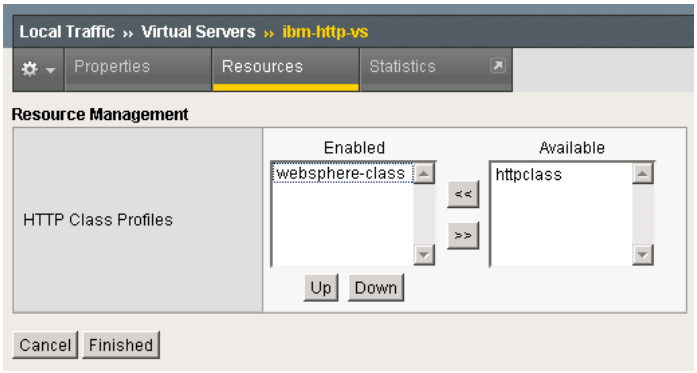

Figure 2.2 *Adding the HTTP Class to the Virtual Server*

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 1-8) 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.*

To create the HTTP profile, use **Creating an HTTP profile**, *on page 1-8,* **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 **IBM WebSphere**.
4. In the **Description** box, optionally type a description for this application.
5. From the **Local Policies** list, select **IBM WebSphere**. This is a pre-defined policy created specifically for IBM WebSphere devices.
6. In the **Requested Host** box, type the host name that your end users use to access the WebSphere deployment. This should be the same host name you used in Step 6a in the HTTP class procedure. In our example, we type **websphere.siterequest.com**.
   If you have additional host names, click the **Add Host** button and enter the host name(s) (see Figure 2.3).
7. Click the **Save** button.

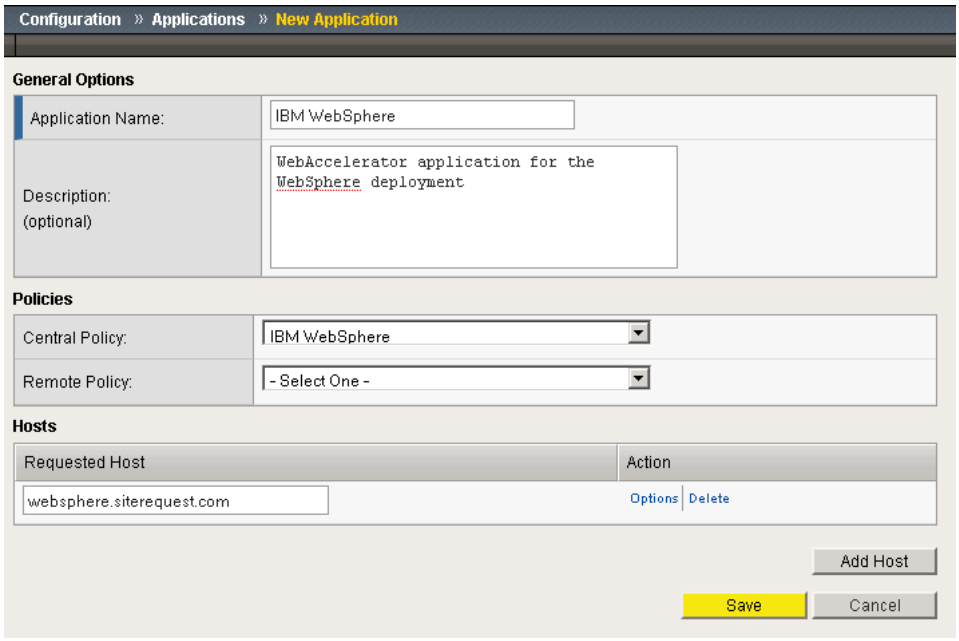

Figure 2.3 *Configuring an Application on the WebAccelerator*

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.

3

Configuring advanced load balancing and health monitoring

- Configuring the BIG-IP LTM for the WebSphere Application Servers
- Creating the IBM WebSphere application server pool
- Creating the LAN optimized TCP profile
- Creating the persistence iRule
- Creating the virtual servers

Deploying the BIG-IP LTM system and IBM WebSphere Servers

In this solution, the BIG-IP LTM implements application monitoring and advanced load balancing capabilities to optimize traffic flows through all application infrastructure. BIG-IP LTM ensures that traffic from front-end web servers is only sent to available application servers. In addition to the round robin and weighting capabilities included in the WebSphere framework, BIG-IP LTM can track server state, and, for example, send traffic to the fastest server, or to the server with the least connections.

In this section the BIG-IP LTM manages application traffic according to the network knowledge it has about the client, web tier and application tier. IBM WebSphere management tools are used to maintain the servers, and all web and application servers included in the deployment should be equally weighted.

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

Prerequisites and configuration notes

The following are prerequisites for this deployment:

- For the scenario described in this chapter, we assume there are three VLANs available in the deployment architecture: one for the BIG-IP virtual servers, one for the presentation tier and one for the application tier. Also, our deployment places all front end web servers are in the presentation tier VLAN and all application servers are in the application tier VLAN.
- We assume you have configured the BIG-IP LTM for the IBM HTTP Server with WebSphere plugin as described in Chapter 1, *Configuring the BIG-IP LTM with IBM HTTP Server and the WebSphere plugin*, on page 1-3. The configuration in Chapter 1 represents the configuration on our web tier VLAN.

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 3.1 shows an example configuration with a redundant pair of BIG-IP devices and a cluster of WebSphere servers. The HTTP servers and WebSphere application servers are configured to communicate with each other using WebSphere tools. In this configuration, we configure an iRule on the BIG-IP LTM system which uses the application's JSESSIONID for persistence.

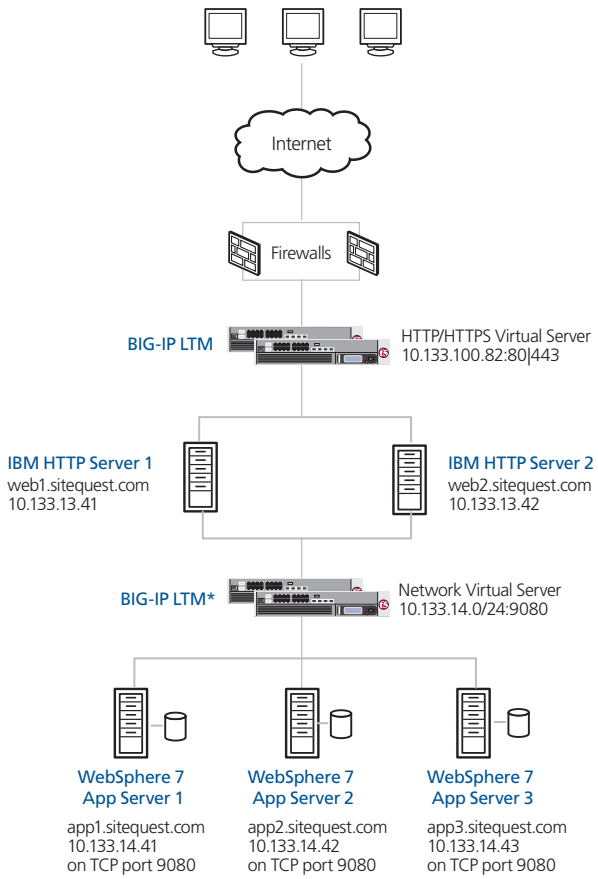

Figure 3.1 *BIG-IP WebSphere configuration example*

** While this is shown as a separate physical BIG-IP system, there may be only one BIG-IP system, with VLANs configured to segment the network.*

Configuring the BIG-IP LTM for the WebSphere Application Servers

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

- *Creating the IBM WebSphere application server pool*
- *Creating the LAN optimized TCP profile*
- *Creating the persistence iRule*
- *Creating the virtual servers*

Tip

We recommend you save your existing BIG-IP configuration before you begin the procedures in this Deployment Guide. For information on backing up or restoring a BIG-IP LTM configuration, refer to the appropriate BIG-IP LTM manual, available on ***Ask F5****.*

Creating the IBM WebSphere application server pool

Use the following procedure to create the WebSphere Application Server pool. This pool uses the health monitor you created in *Creating the HTTP health monitor*, on page 1-4.

To create the pool for the HTTP servers

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. From the Configuration list, select **Advanced**.
4. In the **Name** box, enter a name for your pool.
   In our example, we use **websphere-app-pool**.
5. In the **Health Monitors** section, select the name of the monitor you created in *Creating the HTTP health monitor*, on page 1-4 and click the Add (<<) button. In our example, we select **ibm-web-monitor**.
6. *Optional*: In the **Slow Ramp Time** box, type **300**.
   Because we are using the Least Connections load balancing method, we set the Slow Ramp Time in order to ensure that if a pool member becomes available after maintenance or a new member is added, the BIG-IP does not send all new connections to that member (a newly available member always has the least number of connections).
7. 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 (member)**.

Note: *If you are running multiple applications on the same physical WebSphere application server, we recommend you use* ***Least Connections (Node)****.*

8. For this pool, we leave the Priority Group Activation **Disabled**.
9. In the New Members section, make sure the **New Address** option button is selected.
10. In the **Address** box, add the first server to the pool. In our example, we type **10.133.14.41**.
11. In the **Service Port** box, type the appropriate port for your application. In our example, we type **9080**, the default WebSphere port for the sample applications.
This may be different in your configuration.
12. Click the **Add** button to add the member to the list.
13. Repeat steps 10-12 for each server you want to add to the pool.
14. Click the **Finished** button.

Creating the LAN optimized TCP profile

For this configuration, we create a new LAN optimized TCP profile for the WebSphere application servers. If you 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 **websphere-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 the persistence 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.

To create the iRule

1. On the Main tab, expand **Local Traffic**, click **iRules**, and then click the **Create** button.
2. In the **Name** box, enter a name for your iRule.
   In our example, we use **WebSphereJsessionID**.
3. In the Definition section, type the following iRule (you can also copy and paste, but remove the line numbers).

```
1   when CLIENT_ACCEPTED {
2      set add_persist 1
3   }
4   when HTTP_RESPONSE {
5      if { [HTTP::cookie exists "JSESSIONID"] and $add_persist } {
6         persist add uie [HTTP::cookie "JSESSIONID"]
7         set add_persist 0
8      }
9   }
10 when HTTP_REQUEST {
11     if { [HTTP::cookie exists "JSESSIONID"] } {
12        persist uie [HTTP::cookie "JSESSIONID"]
13     } else {
14       set jsess [findstr [HTTP::uri] "jsessionid" 11 ";"]
15       if { $jsess != "" } {
16           persist uie $jsess
17       }
18     }
19  }
```

Table 3.1 *The JSESSIONID persistence iRule*

4. Click the **Finished** button (see Figure 3.2).

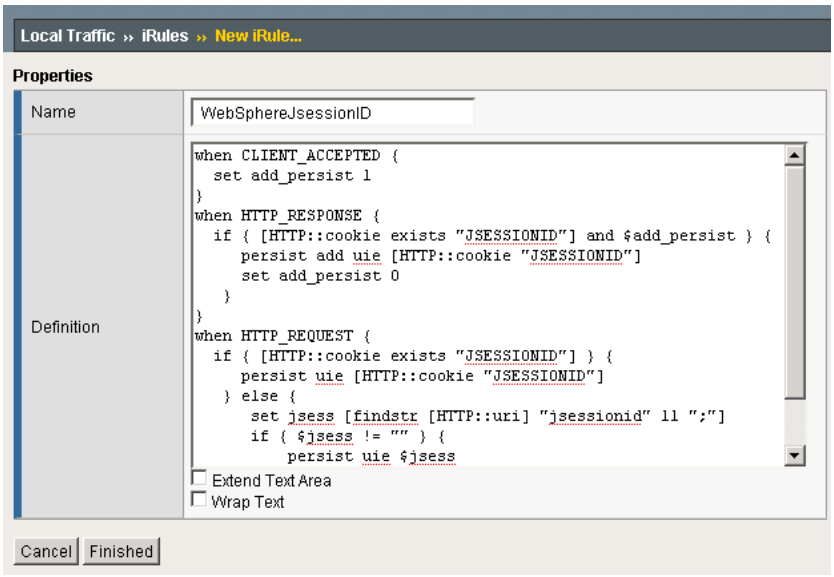

Figure 3.2 *Creating the persistence iRule*

Creating the virtual servers

Next, we configure the virtual servers. For this deployment, we configure three virtual servers, one for the WebSphere application server, and two internal, forwarding virtual servers.

◆ Note

This guide is written with the assumption that you are using VLANs to segment your traffic, and you have VLANs already configured on the BIG-IP LTM for the web tier and the application tier. The forwarding virtual servers you create in this section are locked down specific VLANs. For more information on creating VLANs, see the BIG-IP LTM documentation.

Creating the WebSphere application server virtual server

The first virtual server we create is for the WebSphere application.

To create the application server virtual server

1. On the Main tab, expand **Local Traffic**, click **Virtual Servers** and then click the **Create** button. The New Virtual Server screen opens.
2. In the **Name** box, type a name for this virtual server. In our example, we type **websphere-app-vs**.
3. In the **Destination** section, select the **Network** option button.
4. In the **Address** box, type the IP address of this virtual server. In our example, we use **10.133.14.0**.

5. In the **Netmask** box, type the appropriate netmask. In our example, we type **255.255.255.0**.
6. In the **Service Port** box, type **9080**.

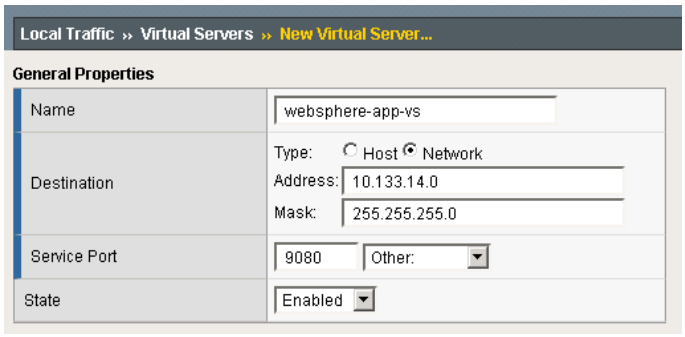

Figure 3.3 *Creating the network virtual server*

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 the LAN optimized TCP profile* section. In our example, we select **websphere-tcp-lan**.
9. From the **HTTP Profile** list, select the name of the profile you created in *Creating an HTTP profile*, on page 1-8. In our example, we select **ibm-http-opt**.
10. In the **VLAN List** row, from the **Available** list, select the VLAN on which your IBM HTTP devices reside, and click the Add (<<) button.
11. In the Resources section, from the **iRules** Available list, select the name of the iRule you created in the *Creating the persistence 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 IBM WebSphere application server pool* section. In our example, we select **websphere-app-pool**.
13. Click the **Finished** button.

Creating the forwarding virtual servers

The next step is to configure the forwarding virtual servers. These virtual servers forward WebSphere management and configuration traffic between the tiers.

To create the forwarding virtual servers

1. On the Main tab, expand **Local Traffic**, click **Virtual Servers**, and then click the **Create** button.
   The New Virtual Server screen opens.
2. In the **Name** box, type a name for this virtual server. In our example, we type **ibm-forward-to-web**.
3. In the **Destination** section, select the **Network** option button.
4. In the **Address** box, type the IP address of this virtual server. In our example, we use **10.133.14.0**.
5. In the **Netmask** box, type the appropriate netmask. In our example, we type **255.255.255.0**.
6. In the **Service Port** box, type **0**, or select ***All Ports** from the list.
7. From the **Configuration** list, select **Advanced**.
8. From the **Type** list, select **Forwarding (IP**).
9. From the **Protocol** list, select ***All Protocols**.
10. In the **VLAN List** row, from the **Available** list, select the VLAN on which your IBM HTTP servers reside, and click the Add (<<) button.
11. Click the **Repeat** button, and repeat this entire procedure, using a unique name such as **ibm-forward-to-web**, the appropriate IP and Netmask (**10.133.100.0** and **255.255.255.0** in our example), and make sure in Step 10 to select the VLAN on which the WebSphere application servers reside. All other settings are the same.
12. Click the Finished button after completing the second virtual server.

This concludes the BIG-IP LTM configuration.

4

Deploying the FirePass controller with IBM WebSphere 7

- Deploying the FirePass controller with IBM WebSphere 7
- Creating groups on the FirePass controller
- Limiting access for the Partner group
- Configuring Endpoint security

Deploying the FirePass controller with IBM WebSphere 7

This section of the deployment guide shows you how to configure F5's FirePass controller for secure remote access to IBM WebSphere deployments. The FirePass controller can integrate seamlessly with IBM to provide authentication, and simple user maintenance.

F5's FirePass® controller is the industry leading SSL VPN solution that enables organizations of any size to provide ubiquitous secure access for employees, partners and customers to applications such as IBM WebSphere, while significantly lowering support costs associated with legacy client-based VPN solutions.

For more information on the FirePass controller, see **http://www.f5.com/products/firepass/**.

Prerequisites and configuration notes

The following are prerequisites for this deployment:

- The FirePass controller should be running version 6.0 or later.
- All of the configuration procedures in this chapter are performed on the FirePass device.
- Our deployment guide configuration uses LDAP authentication, with previously defined LDAP user and group information. Using LDAP authentication is not a requirement for F5 and WebSphere deployments, you can choose a different authentication method. However, the examples in this guide use LDAP.
- This deployment guide is written to the scenario outlined in the following section. It is meant to be a template; modify the configuration as necessary for your deployment.

Configuration scenario

For the scenario used in this deployment guide, the WebSphere deployment, along with an LDAP directory, resides behind a BIG-IP LTM system. There is a requirement to allow employees remote access to all internal resources using the FirePass device. There is also a requirement for trusted partners to access the WebSphere deployment, although only to a limited subset of the portal, with no other access.

This deployment guide describes how to configure the FirePass controller to allow secure remote access to the WebSphere device(s), using LDAP for authentication, and how to configure the FirePass to give one group of users full access, and restrict users in the partner group to a certain directory. This guide also contains procedures on configuring some endpoint security features, including antivirus checks.

Configuring the FirePass controller

To configure the FirePass controller for allowing secure remote access to the IBM WebSphere deployment, you need to complete the following procedures:

- *Connecting to the FirePass controller*
- *Creating groups on the FirePass controller*
- *Limiting access for the Partner group*
- *Configuring Endpoint security*

Connecting to the FirePass controller

To perform the procedures in this Deployment Guide you must have administrative access to the FirePass controller.

To access the Administrative console, in a browser, type the URL of the FirePass controller followed by **/admin/**, and log in with the administrator's user name and password.

Once you are logged on as an administrator, the Device Management screen of the Configuration utility opens. From here, you can configure and monitor the FirePass controller.

Creating groups on the FirePass controller

In this configuration, we configure two types of groups on the FirePass controller, Resource and Master groups. ***Master groups*** contain user information, including details about authentication methods. ***Resource groups*** contain information about applications (resources) that are available to FirePass controller users.

Creating the Resource groups

Resource groups allow you to preconfigure specific applications and access by group, and assign the group to a master group or an individual user. For this configuration, we create two resource groups, one for employees and one for partners, in order to create different access levels to the IBM WebSphere servers.

To configure a resource group

1. From the navigation pane, click **Users**, expand **Groups**, and then click **Resource Groups**.
2. Click the **Create new group** button.
   The Group Management - Create New Group screen opens.
3. In the **New group name** box, type a name for your group and click the **Create** button. In our example we type **employees-WebSphere**.
   The new group appears on the Resource Groups table.

4. From the Resource Groups table, find the row with the name of the group you just created. In this row, from the Portal access column, click **Edit** (see Figure 4.1). The Web Applications section of the Resource Group page opens.

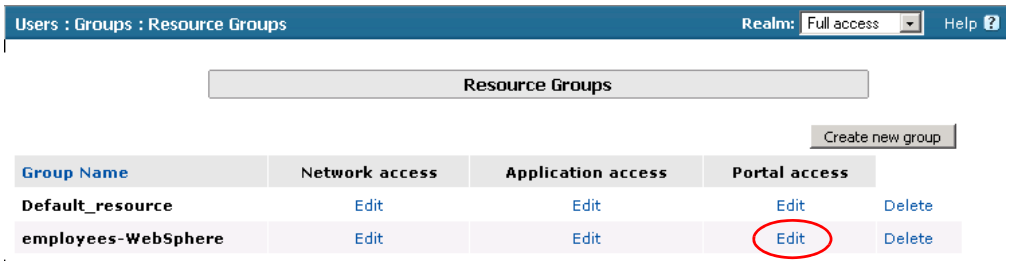

Figure 4.1 *The Resource groups table*

5. Under Web Application Favorites, click **Add New Favorite**. The Favorite options display.
6. Type a name for the Favorite. In our example, we type **WebSphere - Employee**. This Favorite link only displays for members of the Employee group.
7. In the **URL** box, type the URL used to access the IBM WebSphere devices. If you are using a BIG-IP LTM system in front of the IBM WebSphere deployment, this URL should resolve to the IBM HTTP Server virtual server address in DNS. In our example, we type **http://websphere.siterequest.com.**
8. Click the **Add to allow list** link to the right of the URL box. This adds the URL to the list of URLs the users are allowed to access.
9. Configure the rest of the settings as applicable to your deployment.
10. Click the **Add New** button.
The new Favorite is added to the list (see Figure 4.2).

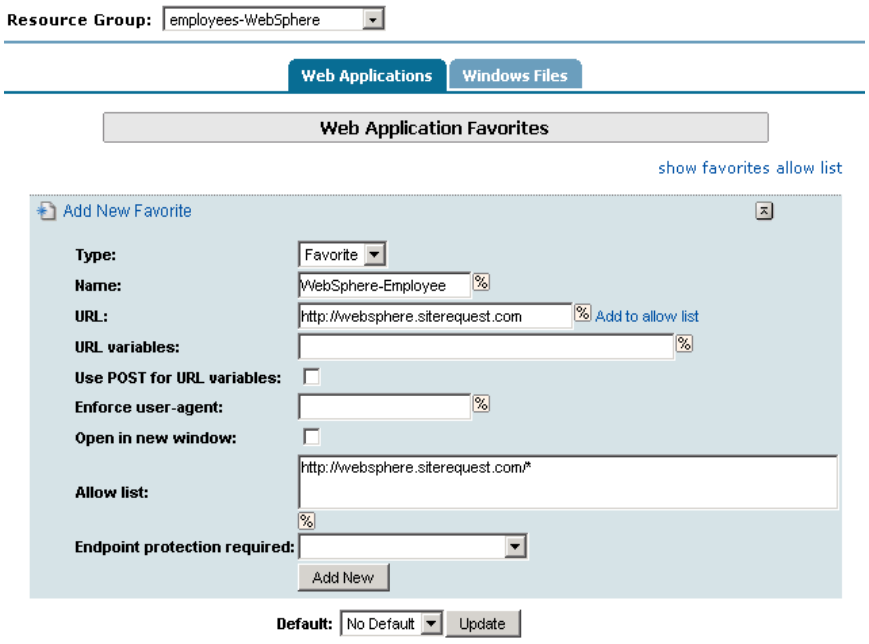

Figure 4.2 *Adding a Web Application Favorite to the Employee group*

11. Repeat this entire procedure for the **Partner** resource group, typing appropriate names for the group and the Favorite.
In our example, we type **partners-WebSphere** for the Resource Group name, and **IBM WebSphere - Partners** for the Favorite name.

    In Step 7, type the path to the appropriate section of the IBM WebSphere deployment that Partners are entitled to access.
For example, the employee Favorite might point to **http://portal.oraclearn.tc.f5net.com**

    while the partner Favorite would point to **http://plantstore.websphere.siterequest.com/partners/**.

Creating the Master groups

FirePass controller master groups are composed of users, authentication methods, and security and policy information. The next task is to create Master groups that will use the resource groups we just created.

To create a new Master Group

1. From the navigation pane, click **Users**, and expand **Groups**.
The Master Groups list screen opens.
2. Click the **Create new group** button.
The Group Management Create New Group screen opens.

3. In the **New group name** box, type the name of your group. In our example we type **WebSphere-LDAP-employee**.
4. In the **Users in group** box, select **External.**
5. From the Authentication method list, select **LDAP**.
6. In the **Copy settings from** list, make sure **Do not copy** is selected (see Figure 4.3).
7. Click the **Create** button.
   The General tab of the new Master Group displays.

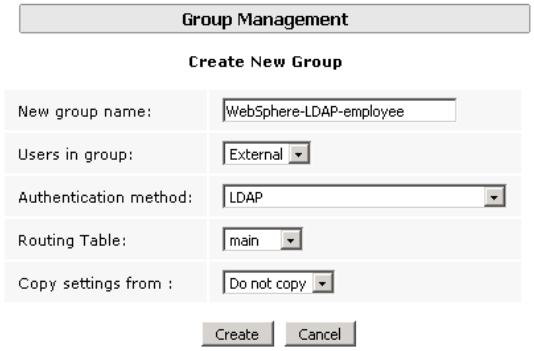

Figure 4.3 *Creating a new Master Group for employees*

8. Click the Resource Groups tab.
   The Resource Groups screen opens.
9. From the **Available** box, select the name of the Resource group you created in the *Creating the Resource groups* section. In our example, we select **employees-WebSphere**.
10. Click the **Add** button to move the group to the **Selected** box, and click the **Update** button. The Resource group is now associated with the Master group.
11. Repeat this procedure to create a new Master group for the Partner resource group. When creating this group, in the **Copy Settings From** list, select the Master group you just created. This way, the authentication settings are automatically created. In our example, we name this group **WebSphere-LDAP-partners**. In Step 9, we select the **partners-WebSphere** Resource group.

Configuring the Master group for LDAP authentication

The next procedure is configuring the Master group to use LDAP authentication.

To configure the Master group to use LDAP authentication

1. From the navigation pane, click **Users**, expand **Groups**, and then click **Master Groups**.

2. Click the name of the Master group you created in the *Creating the Master groups* section. In our example, we select **WebSphere-LDAP-employee**.
3. Click the Authentication tab. The LDAP Authentication tab opens.
4. In the **Host** box, type the host name of your LDAP server. In our example, we type **ldap.siterequest.com**.
5. In the **Port** box, type the port for the LDAP server. In our example we type **636**. If you are not using SSL, the port may be **389**.
6. Check the **Use SSL connection** box (optional).
7. In the **User DN using template** box, type the user DN template. In our example, we type:

   `cn=%logon%,cn=users,dc=siterequest,dc=com`
8. You can optionally click **Test** to test the LDAP authentication.
9. Click the **Save Settings** button.
10. Repeat this entire procedure for the **WebSphere-LDAP-partners** group, using the appropriate settings.

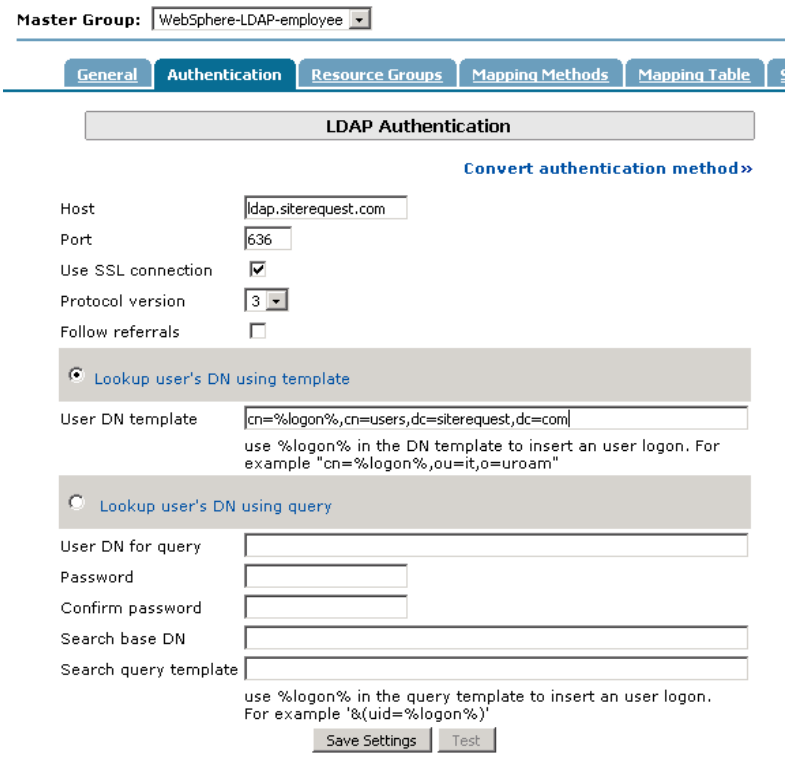

Figure 4.4 *Configuring the Authentication settings*

Limiting access for the Partner group

The FirePass controller allows you to limit access for specific groups on a very granular level. In this scenario, we limit access for the Partner group to only the Favorite we configured earlier, as well as restricting the areas of IBM WebSphere they can access by URL.

To limit access for the Partner group

1. From the navigation pane, click **Portal Access**, and then, under **Web Applications**, click **Master Group Settings**.
2. From the **Master Group** list at the top of the page, select the Master Group you created for the partner group in the *Creating the Master groups* section. In our example, we select **WebSphere-LDAP-partners**.
   The configuration settings for the Master group open.
3. In the **Access limitation** section, make sure there is a check in the **Show administrator-defined favorites only** box.
4. In the **Access Control Lists** section, configure URL pattern matches to allow and deny based on your deployment. In our example, we type the following (separated by commas) in the **Allow** box to restrict the Partner group to these areas of our IBM WebSphere deployment:

   ```
   http://plantstore.websphere.siterequest.com/*
   ```

5. We leave the **Deny** box blank, which allows access to all URLs that pass the allow test (see Figure 4.5). The FirePass checks the deny list, then looks for matches in the allow list, then takes the default action.
   For more information on configuring the Access Control section, see the online help.
6. Leave the **Default Action** list set to **Deny**.
7. Click the **Update** button. The new settings take effect after any users currently logged onto the FirePass controller log out.

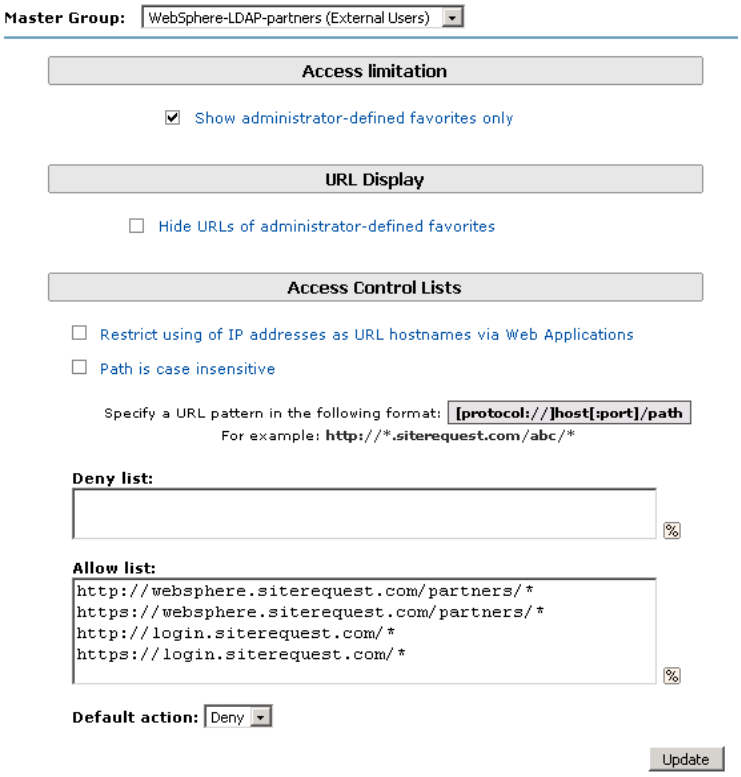

Figure 4.5 *Restricting access to the IBM WebSphere deployment*

Configuring Endpoint security

One of the strong security features of the FirePass controller is the ability to set endpoint security on a extremely granular level.

In the following procedures, we configure a pre-logon check for anti-virus software on Windows machines. The FirePass controller uses this information to deny WebSphere access for members of the Partner Resource group if they do not have the appropriate software. In this configuration example, the FirePass device also denies access to *any* client that is determined to have a virus.

For more information on endpoint security, see the online help.

Creating a pre-logon sequence

The pre-logon sequence allows administrators to create one or more sequences of inspections for items such as installed antivirus programs or OS patch levels.

To configure a pre-logon sequence

1. From the navigation pane, click **Users**, expand **Endpoint Security**, and then click **Pre-Logon Sequence**.
2. In the New Sequence section at the bottom of the page, type a name for the sequence in the **Create New Sequence** box. In our example, we type **WebSphereBasic**.
3. From the **Based on** list, select **template: Collect information with no pre-logon actions**.
4. Click the **Create** button.
   The new sequence appears in the Select Sequence to Use table.
5. ***Warning*** *- Do not click the radio button next to the sequence yet. If you click the radio button, the **Edit** link will be replaced with the **View** link, and you are not able to edit the sequence.*
   In the row of the sequence you just created, click the **Edit** button.
   The Pre-Logon Sequence Editor opens.
6. Move the cursor between **Sequence Start** and the box. A small add **[+]** link appears on the arrow (see the circle marked **1** in Figure 4.6). Click **Add**.
   The Change Sequence panel appears on the right.
7. Click the **Check for Antiviruses** option button, and click the **Apply Changes** button.
   The Edit Action panel opens.

   Note: *The Check for Antiviruses is an optional feature on the FirePass controller. If your device does not have this license, you will not see this option.*
8. Under **Inspectors**, click **Windows Antivirus Checker**.
   The Endpoint Inspector Details page opens in a new window.
9. Configure these options as applicable for your deployment. For more information, click **Help**.
10. Click the **Update** button.
11. In the Sequence pane, find **AV installed**, and click the associated Logon Denied Page link (see the circle marked **2** in Figure 4.6).
    The End Page Properties pane appears on the right.
12. From the Type box, select **Logon Allowed Page**. This allows a user to logon if they have an antivirus checker installed. You can optionally type a message for failed logons.
13. Repeat steps 11 and 12 for the **Fallback** option.
14. ***Optional:*** You can click the Logon Allowed Page or Logon Denied Page links for the other options to produce a custom message when a user is denied access. You can also change the actions taken as a result of the virus checker's findings. For example, you might still want to allow a user to login if there is virus checking software installed, but not currently running.

In our example, we click **Logon Denied Page** next to **Virus Detected**, and type a message informing the user there is a virus on their computer.

15. When you are finished, click **Back to Console** in the upper right corner of the screen (see the circle marked 3 in the following figure).
You return to the Pre-Logon Sequence main page.

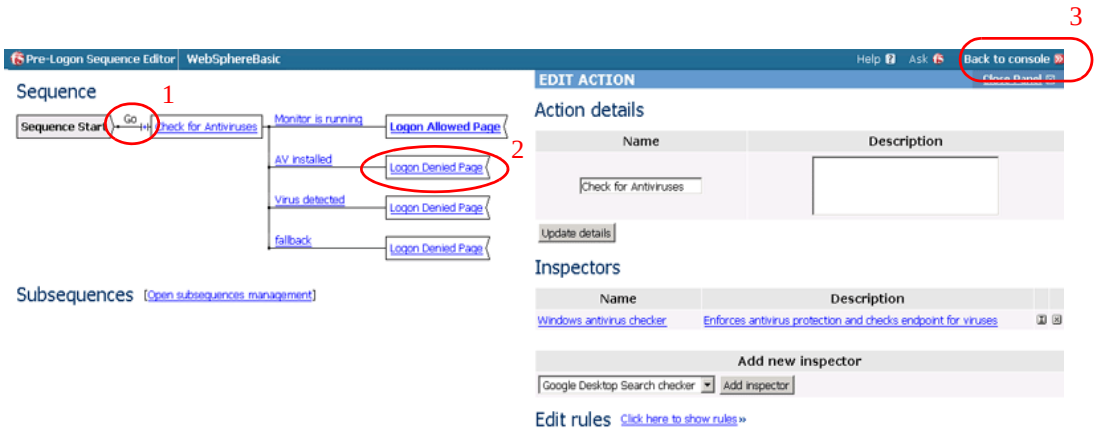

Figure 4.6 *The Pre-Logon Sequence Editor*

16. From the **Select Sequence to Use** section, click the option button next to the sequence you just created. In our example, we click **WebSphereBasic**.
17. Click the **Apply** button.

Protected Configurations

Protected Configurations allow administrators to specify the criteria the endpoint systems must meet to enable access to the various resources**.** In this procedure, we create a protected configuration for the partner group in order make additional security requirements for that group.

To configure Protected Configurations

1. From the navigation pane, click **Users**, expand **Endpoint Security**, and click **Protected Configurations**.
2. Click **New Protection Configuration**.
3. In the Protected configuration ID box, type a name for this configuration. In our example, we type **Partner_config**. You can optionally type a description.

4. Leave the Mode list at the default setting, **Check endpoint protection, grant access if check passed** (see Figure 4.7).

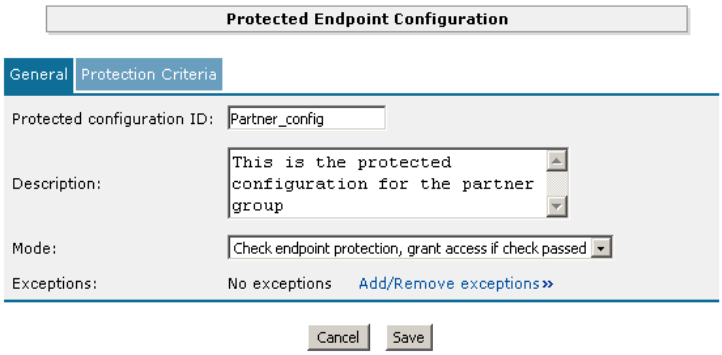

Figure 4.7 *The General tab of the Protected Endpoint Configuration screen*

5. Click the Protected Criteria tab.
6. On the menu bar, click **Information Leaks**.
7. From the Required safety measures or checks list, select **Cache Cleaner** and click the **Add** button. This will remove content from the cache when a user logs off.

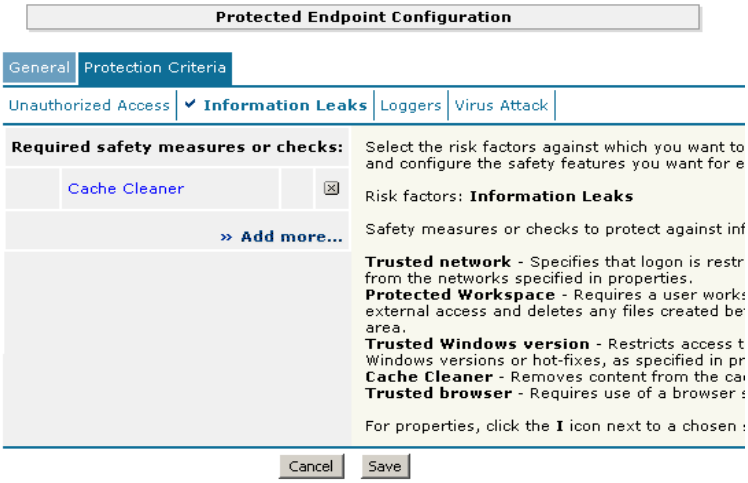

Figure 4.8 *The Protection Criteria tab of the Protected Endpoint Configuration screen*

Important: *The Cache Cleaner feature is currently Windows only. It does not work with Apple Macintosh or Linux systems.*

8. On the menu bar, click **Virus Attack**
9. From the list, select **Antivirus** and click the **Add** button.

10. Click the **I** icon next to Antivirus to configure the antivirus properties (see Figure 4.9). The Select trusted anti-viruses screen opens. Configure these properties as applicable for your configuration, and click the **Save** button.

    You return to the Protection Criteria tab of the Protected Endpoint Configuration page.

11. Click the **Save** button.

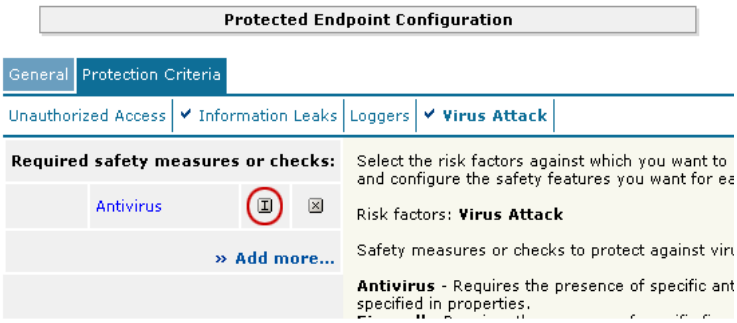

Figure 4.9 *The Edit button for Antivirus properties*

Protecting the Resources

The next step is to associate the protected configuration you just created with a resource.

To protect the resources

1. From the navigation pane, click **Users**, expand **Endpoint Security**, and click **Protect Resources**.
2. From the Resource Table, expand **Web Applications**.
3. Find the **Partners** resource group (in our example, **partners-WebSphere**), and click the **Select** link next to the Favorite you configured.
4. From the **Configuration to protect selected resources** list, select the name of the configuration you created in the preceding procedure. In our example, we select **Partner_config**.
5. From the **Protected Configuration** list, select the name of the configuration you created in the preceding procedure. In our example, we type **Partner_config**.

6. Click the **Select** button.
   A shield image appears in the row.

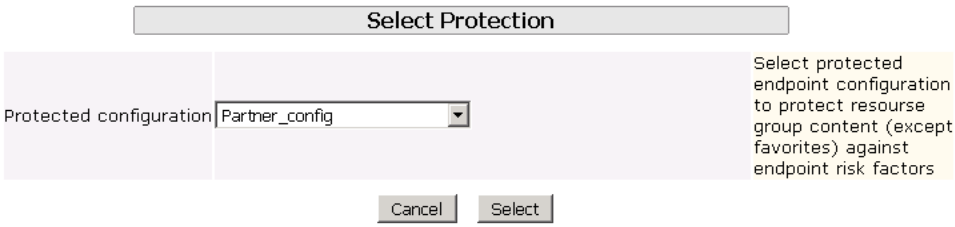

Figure 4.10 *Adding the Protected Configuration to the Resource*

Configuring post-logon actions

The final step is to configure a post-logon action in which the FirePass device injects an Active X control or plug-in to clean the client browser's web cache.

To configure the post-logon action

1. From the navigation pane, click **Users**, expand **Endpoint Security**, and click **Post-Logon Actions**.
2. Click a check in the **Inject ActiveX/Plugin to clean-up client browser web cache** box. A list of options displays.
3. Configure these options as applicable for your deployment. In our example, we leave these options at their default settings.

Conclusion

The FirePass controller is now configured to allow secure remote access to the IBM WebSphere deployment. Remember that the procedures in this Deployment Guide are specific to the scenario described in *Configuration scenario*, on page 4-1. Use this guide as a template, and modify the configuration as applicable to your deployment.