

Deployment Guide

Deploying the BIG-IP LTM System with Microsoft Outlook Web Access



Deploying the BIG-IP LTM System and Microsoft Outlook Web Access

Welcome to the BIG-IP LTM system - Microsoft® Outlook® Web Access Deployment Guide. This guide gives you step-by-step configuration procedures for configuring the BIG-IP LTM system for deployment with Microsoft Outlook Web Access (OWA). It includes procedures on how to direct traffic to the OWA servers, as well using the BIG-IP LTM system for SSL termination.

Microsoft Office Outlook Web Access is an integrated component of Exchange Server 2003. By using only a Web browser and an Internet or intranet connection, Outlook Web Access enables you to read your corporate e-mail messages, schedules, and other information that is stored on a server running Exchange.

For more information on designing a Microsoft Outlook Web Access topology, see

<http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/febetop.mspx>.

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.
- ◆ This Deployment Guide was tested using Exchange Server 2003. While the BIG-IP configuration is the same for Microsoft Exchange Server 2000, it was not tested.
- ◆ All of the configuration procedures in this document are performed on the BIG-IP LTM system. For information on how to deploy or configure Microsoft Outlook Web Access or Microsoft Exchange Server, consult the appropriate Microsoft documentation.
- ◆ Instructions for configuring SSL offloading for Microsoft Exchange Server 2003 and 2000 can be found in the following Microsoft Knowledge base article:
<http://support.microsoft.com/kb/327800/en-us>.

◆ Note

This document is written with the assumption that you are familiar with both the BIG-IP LTM system and Microsoft Exchange and Outlook Web Access. For more information on configuring these products, consult the appropriate documentation.

Configuration example

In this scenario, users connect to a virtual server on the BIG-IP LTM system. If the user initially connects using HTTP, the BIG-IP LTM system redirects the connection to HTTPS using a profile and an iRule. All HTTPS connections are SSL-terminated at the BIG-IP LTM system. The BIG-IP LTM system forwards traffic to available Microsoft Outlook Web Access front-end servers via HTTP, according to specified load balancing and persistence methods. The front-end servers communicate directly with the Microsoft Exchange back-end server(s).

◆ 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.

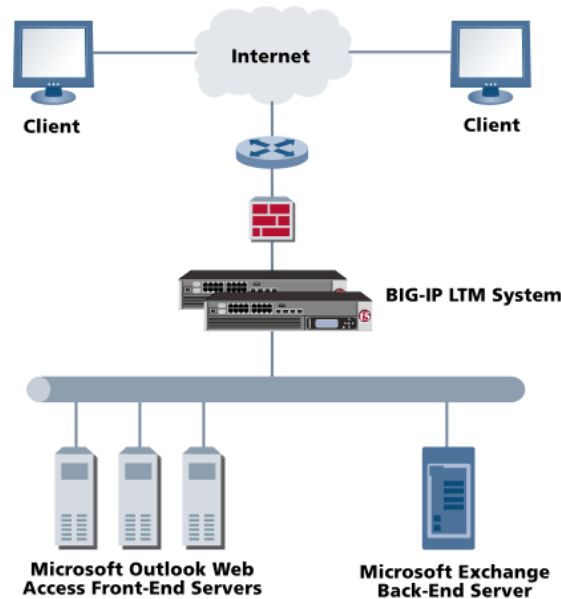


Figure 1 BIG-IP Outlook Web Access configuration example

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 Outlook Web Access

To configure the BIG-IP and SharePoint servers for integration, you need to complete the following procedures:

-
- *Connecting to the BIG-IP device*
 - *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 16.*

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.

Importing keys and certificates

Before you can enable the BIG-IP LTM system to offload SSL traffic from the OWA 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.

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 next task is to create a health monitor for the Microsoft Outlook Web Access devices. This procedure is optional, but very strongly recommended. For this configuration, we use an Extended Content Verification (ECV) 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

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 **OWA_monitor**.

-
- From the **Type** list, select **http**
The HTTP Monitor configuration options appear.
 - 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. In our example, we use a **Interval of 30** and a **Timeout of 91**.
 - Optional:** In the Send String and Receive Rule sections, you can add a Send String and Receive Rule specific to the device being checked.

In our example, we are using the default IIS configuration, so we use a Send String of **iisstart.htm**, and expect the Under Construction page to be returned. If you have modified the IIS configuration on the OWA servers, type a Send String and Receive Rule appropriate for your configuration.

If the page you are requesting in the Send String requires authentication, type a user name and password in the appropriate boxes.

General Properties	
Name	OWA_monitor
Type	HTTP
Import Settings	http
Configuration:	Basic
Interval	30 seconds
Timeout	91 seconds
Send String	GET /iisstart.htm
Receive String	[U]nder [C]onstruction
User Name	
Password	
Reverse	<input type="radio"/> Yes <input checked="" type="radio"/> No
Transparent	<input type="radio"/> Yes <input checked="" type="radio"/> No
Cancel Repeat Finished	

Figure 2 Creating the HTTP monitor

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

Creating the pool

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

To create the OWA pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**. The Pool screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New Pool screen opens.

*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 **OWA_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 **OWA_monitor**.
5. From the **Load Balancing Method** list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network). In our example, we select **Least Connections (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 **10.133.20.11**
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 **80**.
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).

The screenshot displays the configuration utility for a BIG-IP pool. At the top, the configuration is set to 'Basic'. The pool name is 'OWA_pool'. Under 'Health Monitors', 'OWA_monitor' is selected in the 'Active' list, while 'http', 'https', 'https_443', 'siebel_web', and 'tcp' are in the 'Available' list. The 'Resources' section shows 'Least Connections (node)' for the Load Balancing Method and 'Disabled' for Priority Group Activation. In the 'New Members' section, three members are listed: 'R:1 P:1 10.133.20.11 :80', 'R:1 P:1 10.133.20.12 :80', and 'R:1 P:1 10.133.20.13 :80'. The 'Finished' button is highlighted in blue.

Figure 3 Creating the OWA 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.

Creating a cookie persistence profile

The first profile we create is a persistence profile. For this configuration, we create a new cookie persistence profile, based off of the default cookie persistence profile. We use cookie persistence because users must maintain a connection to the same Outlook Web Access device.

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 **OWA_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.
7. Click the **Finished** button.

General Properties	
Name	OWA_cookie
Persistence Type	Cookie
Parent Profile	cookie
Configuration Custom	
Cookie Method	HTTP Cookie Insert <input type="checkbox"/>
Cookie Name	<input type="text"/> <input type="checkbox"/>
Expiration	<input checked="" type="checkbox"/> Session Cookie <input type="checkbox"/>
<input type="button" value="Cancel"/> <input type="button" value="Repeat"/> <input type="button" value="Finished"/>	

Figure 4 Configuring Cookie persistence

Creating an HTTP profile

The next profile we create is an HTTP profile. For this profile, we configure the Header Insert option, where the BIG-IP device inserts the HTTP header FRONT-END-HTTPS: on. For more information on why this header is needed, see <http://tech.f5.com/home/solutions/sol2547.html>.

◆ Note

If you have applied the patch to the Microsoft Exchange Server mentioned in the Microsoft Knowledge Base article 327800 (<http://support.microsoft.com/kb/327800/en-us>) to all of your OWA front-end servers, configuring the Header Insert option on the BIG-IP system may not be necessary. Refer to the article for specific details.

To create a new 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 **OWA_http**.
5. In the Settings section, check the Custom boxes for **Header Insert** and **Redirection Rewrite**.
6. In the Header Insert box, type **FRONT-END-HTTPS: on**.
7. From the **Redirection Rewrite** list, select **All**.
8. Modify any of the other options as applicable for your configuration.
9. Click the **Finished** button.

General Properties	
Name	OWA_http
Parent Profile	http
Settings Custom	
Basic Auth Realm	<input type="checkbox"/>
Fallback Host	<input type="checkbox"/>
Header Insert	FRONT-END-HTTPS: on <input checked="" type="checkbox"/>
Header Erase	<input type="checkbox"/>
Response Chunking	Preserve <input type="checkbox"/>
OneConnect Transformations	<input checked="" type="checkbox"/> Enabled <input type="checkbox"/>
Redirect Rewrite	All <input checked="" type="checkbox"/>
Maximum Header Size	32768 bytes <input type="checkbox"/>
Pipelining	Enabled <input type="checkbox"/>
Insert XForwarded For	Disabled <input type="checkbox"/>
LWS Maximum Columns	80 <input type="checkbox"/>
LWS Separator	<input type="checkbox"/>
Maximum Requests	0 <input type="checkbox"/>

Figure 5 *Creating the HTTP profile*

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

Creating a Client SSL profile

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

To create a new Client SSL profile

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

General Properties	
Name	OWA_clientSSL
Parent Profile	clientssl

Configuration: Basic		Custom
Certificate	OWA	<input checked="" type="checkbox"/>
Key	OWA	<input checked="" type="checkbox"/>

Client Authentication		Custom
Client Certificate	ignore	<input type="checkbox"/>

Cancel Repeat Finished

Figure 6 Creating an SSL profile

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

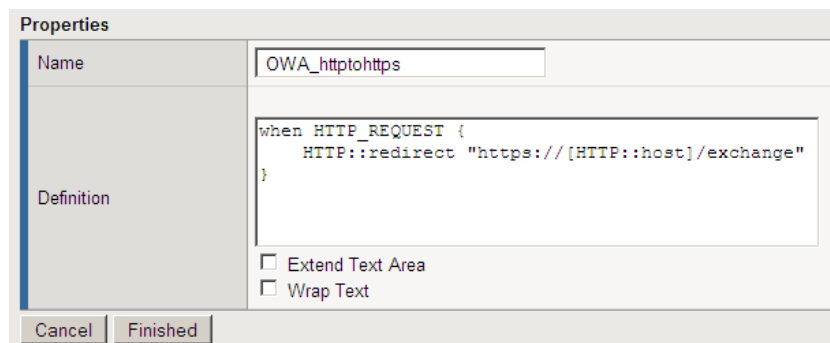
Creating the iRule

The next step is to create an iRule on the BIG-IP LTM system that redirects HTTP traffic to HTTPS. For convenience, the iRule also appends **/exchange** to the URI.

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 **OWA_httpstohttps**.
4. In the Definition section, copy and paste the following iRule:

```
when HTTP_REQUEST {  
    HTTP::redirect "https://[HTTP::host]/exchange"  
}
```



The screenshot shows a 'Properties' dialog box for creating an iRule. The 'Name' field contains 'OWA_httpstohttps'. The 'Definition' field contains the following code:

```
when HTTP_REQUEST {  
    HTTP::redirect "https://[HTTP::host]/exchange"  
}
```

Below the definition field are two checkboxes: 'Extend Text Area' and 'Wrap Text', both of which are unchecked. At the bottom of the dialog are 'Cancel' and 'Finished' buttons.

Figure 7 Creating the iRule

Creating the virtual servers

Next, we configure two virtual servers on the BIG-IP LTM system. The first virtual server is solely to intercept incoming HTTP traffic and redirect it to HTTPS using the iRule you just created. The second virtual server terminates the SSL (HTTPS) connections and forwards traffic via HTTP to the pool of OWA front-end servers.

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 **OWA_http_virtual**.
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 **172.27.92.119**.
6. In the **Service Port** box, type **80**.

General Properties	
Name	<input type="text" value="OWA_http_virtual"/>
Destination	Type: <input checked="" type="radio"/> Host <input type="radio"/> Network Address: <input type="text" value="172.27.92.119"/>
Service Port	<input type="text" value="80"/> <input type="text" value="HTTP"/>
State	<input type="text" value="Enabled"/>

Figure 8 Adding the Outlook Web Access virtual server

7. In the Configuration section, 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 **OWA_http**.
8. In the Resources section, from the **iRule** Available list, select the iRule you created in the *Creating the iRule* section. In our example, we select **OWA_httphttps**.

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

The screenshot shows the 'Resources' configuration window. It is divided into three main sections:

- iRules:** The 'Enabled' list contains 'OWA_httphttps'. The 'Available' list contains '_sys_auth_ldap', '_sys_auth_radius', '_sys_auth_ssl_cc_ldap', '_sys_auth_ssl_ocsp', and '_sys_auth_tacacs'. There are '<<' and '>>' buttons between the lists, and 'Up' and 'Down' buttons below.
- HTTP Class Profiles:** The 'Enabled' list is empty. The 'Available' list contains 'httpclass'. There are '<<' and '>>' buttons between the lists, and 'Up' and 'Down' buttons below.
- Persistence Profiles:** 'Default Pool' is set to 'None'. 'Default Persistence Profile' is set to 'OWA_cookie'. 'Fallback Persistence Profile' is set to 'None'.

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

Figure 9 Resources section of the add virtual server page

- Click the **Finished** button.

Now we create the virtual server for HTTPS.

To create the HTTPS virtual server

- On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
The Virtual Servers screen opens.
- In the upper right portion of the screen, click the **Create** button.
The New Virtual Server screen opens.
- In the **Name** box, type a name for this virtual server. In our example, we type **OWA_https_virtual**.
- In the **Destination** section, select the **Host** option button.
- In the **Address** box, type the IP address of this virtual server. In our example, we use **172.27.92.119**.

- In the **Service Port** box, type **443**, or select HTTPS from the list.

General Properties	
Name	OWA_https_virtual
Destination	Type: <input checked="" type="radio"/> Host <input type="radio"/> Network
	Address: 172.27.92.119
Service Port	443 HTTPS
State	Enabled

Figure 10 Adding the HTTPS virtual server

- In the Configuration section, from the **HTTP Profile** list, select the profile you created in the *Creating an HTTP profile* section. In our example, we select **OWA_http**.
- 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 **OWA_clientSSL**.
- In the Resources section, from the **Default Pool** list, select the pool you created in the *Creating the pool* section. In our example, we select **OWA_pool**.
- From the **Default Persistence Profile** list, select the persistence profile you created in the *Creating a cookie persistence profile* section. In our example, we select **OWA_cookie**.

Resources	
iRules	<div style="display: flex; justify-content: space-between;"> <div>Enabled</div> <div>Available</div> </div> <div style="display: flex; align-items: center;"> <div style="border: 1px solid gray; width: 150px; height: 40px; margin-right: 5px;"></div> <div style="margin: 0 5px;"> << << >> >> </div> <div style="border: 1px solid gray; width: 150px; height: 40px; padding: 2px;"> _sys_auth_ldap _sys_auth_radius _sys_auth_ssl_cc_ldap _sys_auth_ssl_ocsp _sys_auth_tacacs </div> </div> <div style="display: flex; justify-content: center; margin-top: 5px;"> Up Down </div>
HTTP Class Profiles	<div style="display: flex; justify-content: space-between;"> <div>Enabled</div> <div>Available</div> </div> <div style="display: flex; align-items: center;"> <div style="border: 1px solid gray; width: 150px; height: 40px; margin-right: 5px;"></div> <div style="margin: 0 5px;"> << << >> >> </div> <div style="border: 1px solid gray; width: 100px; height: 40px; padding: 2px;"> httpclass </div> </div> <div style="display: flex; justify-content: center; margin-top: 5px;"> Up Down </div>
Default Pool	<div style="display: flex; align-items: center;"> + <div style="border: 1px solid gray; padding: 2px;">OWA_pool</div> </div>
Default Persistence Profile	<div style="border: 1px solid gray; padding: 2px;">OWA_cookie</div>
Fallback Persistence Profile	<div style="border: 1px solid gray; padding: 2px;">None</div>
<div style="display: flex; justify-content: center; gap: 10px;"> Cancel Repeat Finished </div>	

Figure 11 Resources section of the add virtual server page

-
11. Click the **Finished** button.

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_OWA_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.