



Deploying the BIG-IP System v11 with HTTP Applications

What's inside:

- 2 Prerequisites and configuration notes
- 2 Configuration example
- 3 Preparation Worksheet
- 4 Configuring the BIG-IP iApp for HTTP applications
- 8 Next steps
- 10 Appendix: Manual configuration table
- 12 Document Revision History

Welcome to the F5 deployment guide for HTTP applications. This document contains guidance on configuring the BIG-IP system version 11 for simple web server implementations, resulting in a secure, fast, and available deployment.

BIG-IP version 11.0 introduces iApp™ Application templates, an extremely easy way to configure the BIG-IP system for your HTTP applications.

Why F5?

The BIG-IP system provides a number of ways to accelerate, optimize, and scale HTTP deployments. When BIG-IP LTM relieves web servers from tasks such as compression, caching, and SSL processing, each server is able to devote more resources to running applications and can service more user requests.

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

Products and versions tested

Product	Version
BIG-IP LTM	v11

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

What is F5 iApp™?

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

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

Document Version

1.1


Prerequisites and configuration notes

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

- For this deployment guide, the BIG-IP LTM system **must** be running version 11.0 or later. If you are using a previous version of the BIG-IP LTM system, see the Deployment Guide index on F5.com. The configuration described in this guide does not apply to previous versions.
- This deployment guide provides guidance for using the iApp for HTTP applications found in version 11.0 and later. For advanced users extremely familiar with the BIG-IP, there is a manual configuration table at the end of this guide. However, we strongly recommend using the iApp template.
- If you are using the BIG-IP system to offload SSL, we assume you have already obtained an SSL certificate and key, and it is installed on the BIG-IP LTM system.
- This document contains guidance on optional modules, including Application Visibility Reporting and WebAccelerator. To take advantage of these modules, they must be licensed and provisioned before starting the iApp template. For more information on licensing modules, contact your sales representative. Note that AVR is licensed on all systems, but must be provisioned before beginning the iApp template.
- If you are deploying a specific application, we recommend you first check the deployment guide index at <http://www.f5.com/solutions/resources/deployment-guides.html> to see if there is a deployment guide for your specific application.

Configuration example

In this Deployment Guide, the BIG-IP system is optimally configured to optimize and direct traffic to HTTP servers. The following diagram shows a logical configuration example with a redundant pair of BIG-IP LTM devices running the optional WebAccelerator module, in front of a group of HTTP servers.



Preparation Worksheet

In order to use the iApp for HTTP applications, you need to gather some information, such as server IP addresses and domain information. Use the following worksheet to gather the information you will need while running the template. The worksheet does not contain every question in the template, but rather includes the information that is helpful to have in advance. More information on specific template questions can be found on the individual pages.

You might find it useful to print this table and then enter the information.

➡ **Note:** *Although we show space for 10 pool members, you may have more or fewer members in each pool.*

IP Addresses/FQDN	SSL Offload	Pool Members	Sync/Failover Groups*	TCP request queuing*	WAN or LAN clients
<p>IP address you will use for the LTM virtual server:</p> <p>FQDN that will resolve to the virtual server address:</p>	<p><i>Offloading SSL?</i> Yes No</p> <p>If offloading SSL, import a certificate and key into the BIG-IP LTM before running the template.</p> <p>Certificate:</p> <p>Key:</p>	<p>HTTP server IP addresses:</p> <p>1:</p> <p>2:</p> <p>3:</p> <p>4:</p> <p>5:</p> <p>6:</p> <p>7:</p> <p>8:</p> <p>9:</p> <p>10:</p> <p>Port used by the server:</p>	<p>If using the Advanced feature of Sync/Failover Groups, you must already have a Device Group and a Traffic Group</p> <p>Device Group name:</p> <p>Traffic Group name:</p>	<p>If using TCP request queuing, you should know the queue length and timeout, as well as the connection limit for the node.</p> <p>Request queue length:</p> <p>Timeout:</p> <p>Node Connection limit:</p>	<p>Most clients connecting through BIG-IP to the HTTP server are coming over a:</p> <p>LAN</p> <p>WAN</p>
<p>Optional Modules (you must have provisioned modules before running the template)</p>					
<p><i>Application Visibility Reporting (AVR)*</i></p>		<p><i>WebAccelerator*</i></p>			
<p>If using AVR, we strongly recommend you first create a custom Analytics profile before running the template.</p> <p>Analytics profile name:</p>		<p>All FQDNs for the HTTP application:</p> <p>1:</p> <p>2:</p> <p>3:</p> <p>4:</p> <p>5:</p>			

* *Optional*

Configuring the BIG-IP iApp for HTTP applications

Use the following guidance to help configure the BIG-IP system for HTTP applications using the BIG-IP iApp template.

Getting Started with the iApp for HTTP applications

To begin the HTTP iApp Template, use the following procedure.

1. Log on to the BIG-IP system.
2. On the Main tab, expand **iApp**, and then click **Application Services**.
3. Click **Create**. The Template Selection page opens.
4. In the **Name** box, type a name. In our example, we use **HTTP-app_**.
5. From the **Template** list, select **f5.http**.
The HTTP template opens.

Advanced options

If you select Advanced from the Template Selection list, you see Sync and Failover options for the application. This feature, new to v11, is a part of the Device Management configuration. This functionality extends the existing High Availability infrastructure and allows for clustering, granular control of configuration synchronization and granular control of failover. For more information on Device Management, see the Online Help or product documentation.

1. **Configure Sync/Failover?**
If you want to configure the Application for Sync or failover groups, select **Yes** from the list.
 - a. **Device Group**
If you select Yes from the question above, the Device Group and Traffic Group options appear. If necessary, uncheck the Device Group box and then select the appropriate Device Group from the list.
 - b. **Traffic Group**
If necessary, uncheck the Traffic Group box and then select the appropriate Traffic Group from the list.

Analytics

The first section of the template asks questions about Analytics. The Application Visibility Reporting (AVR) module allows you to view statistics specific to your HTTP implementation. AVR is available on all BIG-IP systems v11 and later, however you must have the AVR provisioned for this option to appear. Note that these are only for application visibility reporting, you can view object-level statistics from the BIG-IP without provisioning AVR.

If you plan on using AVR for analytics, we recommend creating a custom Analytics profile before beginning the template. To create a new profile, from the Main tab, select Profiles and then click Analytics. Click New and then configure the profile as applicable for your configuration. See the online help or product documentation for specific instructions.

1. **Enable Analytics**
Choose whether you want to enable AVR for Analytics.

Tip

If using AVR, create a new Analytics profile before beginning the iApp for more specific reporting

2. **Analytics Profile**

You must decide whether to use the default Analytics profile, or create a new one. As mentioned previously, we recommend creating a new profile to get the most flexibility and functionality out of AVR. If you choose to create a new profile after starting the template, you must exit the template, create the profile, and then restart the template.

To use the default Analytics profile, choose Use **Default Profile** from the list.

To choose a custom profile, leave the list set to **Select a Custom Profile**, and then from the Analytics profile list, select the custom profile you created.

Virtual Server Questions

The next section of the template asks questions about the BIG-IP virtual server. A virtual server is a traffic-management object on the BIG-IP system that is represented by an IP address and a service. Clients can send application traffic to a virtual server, which then directs the traffic according to your configuration instructions.

1. **IP address for the virtual server**

This is the address clients use to access this application (or a FQDN will resolve to this address). You need an available IP address to use here.

2. **Port**

Type the port you want to use for this virtual server.

Important



If you are using the BIG-IP LTM to offload SSL processing from the servers (as described in *SSL Encryption questions on page 6*), the system automatically creates a virtual server on port 443, with an additional virtual server on the port you specify in this Port box (port 80 by default) which redirects users to the port 443 virtual server. **If you are offloading SSL to the BIG-IP system, do NOT type 443 in this box.**

If you are not offloading SSL on the BIG-IP system, type 443 or the appropriate port here.

3. **Routes or secure network address translation**

If the servers do not have a route back for clients through the BIG-IP, i.e. if they do not use the BIG-IP as the default gateway, the BIG-IP uses Secure Network Address Translation (SNAT) Automap (exception in #4) to translate the client's source address to an address configured on the BIG-IP. The servers then use this new source address as the destination address when responding to traffic originating through the BIG-IP.

If you indicate the servers do have a route back to the clients through the BIG-IP, the BIG-IP does not translate the client's source address; in this case, you must make sure the BIG-IP is configured as the gateway to the client networks (usually the default gateway) on the servers.

We recommend choosing **No** from the list because it is secure and does not require you to configure routing manually.

If you are configuring your BIG-IP LTM in a "one-armed" configuration with your servers -- where the BIG-IP virtual server(s) and the servers have IP addresses on the same subnet -- you must choose No.

4. **More than 64,000 simultaneous connections**

If you do not expect more than 64,000 simultaneous connections, leave this answer set to

No and continue with #5.

If you have a large deployment and expect more than 64,000 connections at one time, the iApp creates a SNAT Pool instead of using SNAT Automap. With a SNAT Pool, you need one IP address for each 64,000 connections you expect. Select **Yes** from the list. A new row appears with an IP address field. In the **Address** box, type an IP address and then click **Add**. Repeat for any additional IP addresses.

5. **NTLM**

If you have configured the servers to use NTLM authentication, select Yes from the list. If the servers do not use NTLM, leave the list set to No.

SSL Encryption questions

Before running the template you should have already imported a certificate and key onto the BIG-IP system. While the BIG-IP system does include a self-signed SSL certificate that can be used internally or for testing, we strongly recommend importing a certificate and key issued from a trusted Certificate Authority.

For information on SSL certificates on the BIG-IP system, see the online help or the *Managing SSL Certificates for Local Traffic* chapter in the **Configuration Guide for BIG-IP Local Traffic Manager** available at <http://support.f5.com/kb/en-us.html>.

To configure the BIG-IP to offload SSL, select **Yes** from the list.

1. **Certificate**

Select the certificate for you imported for the application from the certificate list.

2. **Key**

Select the associated key from the list.

Server Pool, Load Balancing, and Service Monitor questions

In this section, you add the HTTP servers, and configure the health monitor and pool.

1. **New Pool**

Choose **Create New Pool** unless you have already made a pool on the LTM for the HTTP devices.

2. **Load balancing method**

While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (member)**.

3. **Address/Port**

Type the IP Address and Port for each server. You can optionally add a Connection Limit (see note on the left). Click **Add** to add additional servers to the pool.

4. **TCP Request Queuing**

TCP request queuing provides the ability to queue connection requests that exceed the capacity of connections for a pool as determined by the connection limit. Consequently, instead of dropping connection requests that exceed the capacity of a pool, TCP request queuing enables those connection requests to reside within a queue in accordance with defined conditions until capacity becomes available. For more information on TCP Request Queuing, see the *New Features Guide for BIG-IP Version 11*, available on Ask F5.

If you want the BIG-IP to queue TCP requests, select **Yes** from the list. Additional options appear.

a. Type a queue length in the box. Leave the default of 0 for unlimited.

🔔 **Important**

If you enable TCP Request Queuing, you must have a Connection Limit set on at least one of the nodes when configuring the Address/Port

- b. Type a number of milliseconds for the timeout value.
5. **Health Monitor**
Choose **Create New Monitor** unless you have already made a health monitor on the LTM for the HTTP application.
6. **Interval**
Specifies how often the system checks the health of the servers. We recommend the default of 30 seconds.
7. **HTTP Request**
This is optional. You can configure the template to retrieve a specific page by typing the path here. Leaving the default (GET /) marks the node up if anything is returned from the web page.
8. **HTTP version**
Unless the majority of your users are using HTTP 1.0, we recommend selecting **Version 1.1** from the list.
 - FQDN: When you select Version 1.1, a new row appears asking for the FQDN the clients use to access the application. Type it here.
9. **Monitor response string**
Optional. If you configured a unique HTTP Request, this is where you enter the expected response.

Protocol Optimization Questions

In this section, you configure protocol optimizations.

1. **WAN or LAN**
Specify whether most clients are connecting over a WAN or LAN.
2. **WebAccelerator**
If you have licensed and provisioned the WebAccelerator module, you have the option of using it for the HTTP application. The WebAccelerator provides application acceleration for remote users.
 - a. *DNS names*
If you select Yes, an additional row appears in the template asking for the fully qualified domain names used for this application. The BIG-IP system uses these entries for the Requested Hosts field, allowing the WebAccelerator module to accelerate the traffic to these virtual hosts.

In the **Host** box, type the **FQDN**. If you have additional FQDNs, click the **Add** button.
 - b. *X-WA-info Header*
By default, the WebAccelerator X-WA-info header is not included in the response from the BIG-IP. This header is useful for debugging WebAccelerator behavior. There are two additional options:
 - Standard: If you choose Standard, the BIG-IP inserts a HTTP header that includes numeric codes which indicate if and how each object was cached.
 - Debug: If you choose Debug, the BIG-IP includes extended information which may help for extended troubleshooting.
 - c. *WebAccelerator Performance monitor*
While the BIG-IP Dashboard provides statistics and performance graphs related to WebAccelerator, you can choose to enable the WebAccelerator performance monitor for

legacy WebAccelerator performance monitoring for debugging purposes. The results can be found in the Main tab of the navigation page, under WebAccelerator, by clicking Traffic Reports.

In our example, we leave the performance monitor **Disabled**.

d. *WebAccelerator policy*

For this template, F5 recommends the **Generic Policy - Enhanced** policy to achieve the best results for Web acceleration of HTTP traffic. Should F5 publish an updated policy to DevCentral that you have downloaded and imported, or if a custom policy is created for your environment (locally), you can select that custom policy from the list. In our example, we leave the default.

Finished

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

Next steps

After completing the iApp Template, the BIG-IP Application Services page opens for the HTTP service you just created. To see the list of all the configuration objects created to support the HTTP application, on the Menu bar, click **Components**. The complete list of all related objects opens. You can click individual objects to see the settings.

Once the objects have been created, you are ready to use the new deployment.

Modifying DNS settings to use the BIG-IP virtual server address

Before sending traffic to the BIG-IP system, your DNS administrator may need to modify any DNS entries for the HTTP implementation to point to the BIG-IP system's virtual server address.

Modifying the iApp configuration

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

To modify the configuration

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

Viewing statistics

You can easily view a number of different statistics on the BIG-IP system related to the configuration objects created by the iApp template. You can get statistics specific to the Application Service if

you have provisioned AVR. Otherwise, you can always get object-level statistics.

AVR statistics

If you have provisioned AVR, you can get application-level statistics for your HTTP application service.

To view AVR statistics

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. From the Application Service List, click the HTTP service you just created.
3. On the Menu bar, click **Analytics**.
4. Use the tabs and the Menu bar to view different statistics for your iApp.

Object-level statistics

If you haven't provisioned AVR, or want to view object-level statistics, use the following procedure.

To view object-level statics

1. On the Main tab, expand **Overview**, and then click **Statistics**.
2. From the **Statistics Type** menu, you can select **Virtual Servers** to see statistics related to the virtual servers.
3. You can also choose **Pools** or **Nodes** to get a closer look at the traffic.
4. To see Networking statistics in a graphical format, click **Dashboard**.

For more information on viewing statistics on the BIG-IP system, see the online help or product documentation.

Appendix: Manual configuration table

We strongly recommend using the iApp template to configure the BIG-IP LTM for HTTP applications. Advanced users extremely familiar with the BIG-IP system can use the following table to manually configure the BIG-IP system.

The following table contains a list of BIG-IP LTM configuration objects along with any non-default settings you should configure as a part of this deployment. Unless otherwise specified, settings not mentioned in the table can be configured as applicable for your configuration. For specific instructions on configuring individual objects, see the online help or product manuals.

BIG-IP LTM Object	Non-default settings/Notes		
Health Monitor (Main tab-->Local Traffic -->Monitors)	Name	Type a unique name	
	Type	HTTP	
	Interval	30 (recommended)	
	Timeout	91 (recommended)	
Pool (Main tab-->Local Traffic -->Pools)	Name	Type a unique name	
	Health Monitor	Select the monitor you created above	
	Slow Ramp Time¹	300	
	Load Balancing Method	Choose a load balancing method. We recommend Least Connections (Member)	
	Address	Type the IP Address of the nodes	
	Service Port	80 (click Add to repeat Address and Service Port for all nodes)	
Profiles (Main tab-->Local Traffic -->Profiles)	HTTP (Profiles-->Services)	Name Parent Profile Rewrite Redirect?	Type a unique name http Matching
	TCP WAN (Profiles-->Protocol)	Name Parent Profile	Type a unique name tcp-wan-optimized
	TCP LAN (Profiles-->Protocol)	Name Parent Profile	Type a unique name tcp-lan-optimized
	Persistence (Profiles-->Persistence)	Name Persistence Type	Type a unique name Cookie
	OneConnect (Profiles-->Other)	Name Parent Profile	Type a unique name oneconnect
	Client SSL² (Profiles-->SSL)	Name Parent Profile Certificate and Key	Type a unique name clientssl Select the Certificate and Key you imported from the associated list
	Web Acceleration (Profiles-->Services)	Name Parent Profile	Type a unique name optimized-caching
		Name Parent Profile	Type a unique name wan-optimized-compression
			application/vnd.ms-publisher
			application/(xls excel msexcel ms-excel x-excel x-xls x-msexcel x-ms-excel vnd.excel vnd.msexcel vnd.ms-excel)
			application/(word doc msword winword ms-word x-word x-msword vnd.word vnd.msword vnd.ms-word)
	HTTP Compression (Profiles-->Services)	Content List -->Include List (Add each entry to the Content Type box and then click Include)	application/(xml x-javascript javascript x-ecmascript ecmascript) application/(powerpoint mspowerpoint ms-powerpoint x-powerpoint x-mspowerpoint vnd.powerpoint vnd.ms-powerpoint vnd.ms-powerpoint vnd.ms-pps) application/(mpp msproject x-msproject x-ms-project vnd.ms-project) application/(visio x-visio vnd.visio vsd x-vsd x-vsd) application/(pdf x-pdf acrobat vnd.pdf)

¹ You must select **Advanced** from the **Configuration** list for these options to appear

² Only required if offloading SSL on the BIG-IP LTM

Configuration table, continued

BIG-IP LTM Object	Non-default settings/Notes	
Virtual Servers (Main tab-->Local Traffic -->Virtual Servers)	HTTP	
	Name	Type a unique name.
	Address	Type the IP Address for the virtual server
	Service Port	80
	Protocol Profile (client)^{1,2}	Select the WAN optimized TCP profile you created above
	Protocol Profile (server)^{1,2}	Select the LAN optimized TCP profile you created above
	HTTP Profile²	Select the HTTP profile you created above
	Web Acceleration profile²	Select the Web Acceleration profile you created above
	HTTP Compression profile²	Select the HTTP Compression profile you created above
	OneConnect²	Select the OneConnect profile you created above
	SNAT Pool³	Automap (optional; see footnote ³)
	Default Pool²	Select the pool you created above
	Persistence Profile²	Select the Persistence profile you created
	iRule⁴	If offloading SSL only: Enable the built-in _sys_https_redirect irule
	HTTPS⁵	
	Name	Type a unique name.
	Address	Type the IP Address for the virtual server
	Service Port	443
	Protocol Profile (client)¹	Select the WAN optimized TCP profile you created above
	Protocol Profile (server)¹	Select the LAN optimized TCP profile you created above
	HTTP Profile	Select the HTTP profile you created above
	Web Acceleration profile	Select the Web Acceleration profile you created above
	HTTP Compression profile	Select the HTTP Compression profile you created above
OneConnect	Select the OneConnect profile you created above	
SSL Profile (client)	Select the Client SSL profile you created above	
SNAT Pool²	Automap (optional; see footnote ³)	
Default Pool	Select the pool you created above	
Persistence Profile	Select the Persistence profile you created	

¹ You must select **Advanced** from the **Configuration** list for these options to appear

² Do not enable these objects on the HTTP virtual server if offloading SSL. The HTTP virtual server is only used for redirecting users to the HTTPS virtual server, and only requires a name, IP address, Port, and the redirect iRule.

³ If want to use SNAT, and you have a large deployment expecting more than 64,000 simultaneous connections, you must configure a SNAT Pool with an IP address for each 64,000 simultaneous connections you expect. See the BIG-IP documentation on configuring SNAT Pools.

⁴ Only enable this iRule if offloading SSL

⁵ Only create this virtual server if offloading SSL

Document Revision History

Version	Description
1.0	New Version
1.1	Clarified guidance on the service port to use on the BIG-IP virtual server.

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

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

F5 Networks
Asia-Pacific
apacinfo@f5.com

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

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

