



Deploying the BIG-IP LTM with Multiple BIG-IP WebAccelerator Devices

What's inside:

- 2 Prerequisites and configuration notes
- 2 Configuration example
- 3 Configuring the BIG-IP LTM interior virtual server
- 6 Configuring the WebAccelerator devices
- 7 Configuring the BIG-IP LTM exterior virtual server
- 8 Troubleshooting
- 9 Appendix: Optional EAV monitor based on CPU usage

Welcome to the F5 Deployment Guide for the F5 BIG-IP® Local Traffic Manager™ (LTM) with multiple BIG-IP WebAccelerator devices. This guide shows you how to configure the BIG-IP LTM and multiple WebAccelerator devices for fast and reliable access to your applications.

This document is written for organizations with heavy traffic loads who require more than a single WebAccelerator device for their application deployment. The BIG-IP LTM intelligently directs traffic to a pool of WebAccelerator devices, which accelerates the traffic between the application and the end user. The BIG-IP uses sophisticated health monitors not only to ensure that traffic is directed to available WebAccelerator devices, but also to provide intelligent traffic management based on the utilization of the WebAccelerator devices, resulting in the best possible user experience.

The BIG-IP WebAccelerator provides a series of intelligent technologies that overcome performance issues involving browsers, web application platforms, and WAN latency. By decreasing page download times, BIG-IP WebAccelerator offloads servers, decreases bandwidth usage, increases revenue, and ensures the productivity of application end users.

For more information on the F5 BIG-IP system, see <http://www.f5.com/products/big-ip/>

Products and versions tested

Product	Version
BIG-IP LTM and WebAccelerator	10.2.1, 10.2.2 (applies to versions 10.x)

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

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

Prerequisites and configuration notes

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

- You must be running BIG-IP version 10.x. The configuration in this guide does not apply to BIG-IP version 11.0 or later.
- For the configuration in this guide, you should have at least two active WebAccelerator devices (and not just an active/standby high availability pair).
- The BIG-IP system must be initially configured with the proper VLANs and Self IP addresses. For more information on VLANs and Self IPs, see the online help or the BIG-IP documentation.

Configuration example

In the configuration described in this guide, a client requests a web application. The exterior virtual server on the BIG-IP LTM receives the request and intelligently directs the request to an available WebAccelerator in a pool of WebAccelerator devices. The WebAccelerator device uses an acceleration policy to optimize the transaction, and then sends the request to the interior LTM virtual server. The LTM then intelligently directs the request to the best available web application server.

You can host both the internal and external virtual servers on the same BIG-IP LTM, or you may have a separate internal and external BIG-IP LTM devices. In the following logical configuration example, we show separate BIG-IP LTM devices for clarity.

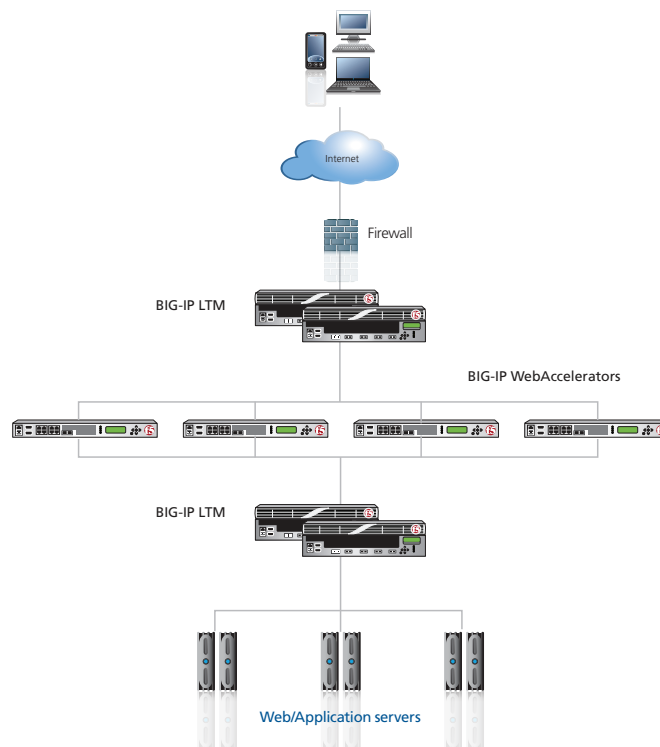
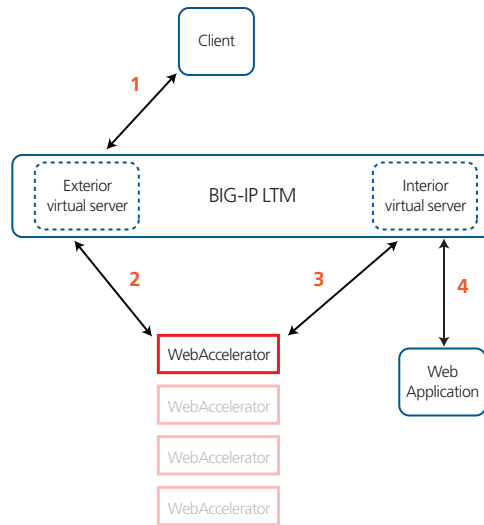


Figure 1: Configuration example

The following shows the traffic flow in a configuration using a single BIG-IP LTM.



Configuring the BIG-IP LTM interior virtual server

In this section, we configure the interior virtual server on the BIG-IP LTM. As mentioned previously, this virtual server can be on the same physical device as the exterior virtual server, or on separate devices.


The interior virtual server is for your web application. In the following procedures, we use a generic HTTP web application as an example. You can modify the BIG-IP configuration objects, such as the health monitor and the profiles, to suit your particular application.

Configuring the VLANs

Configuring VLANs and Self IP address should already be complete before beginning the configuration in this guide. This section describes our VLAN configuration as a reference for possible deployments.

In our example, we have a total of three VLANs (and associated self IP addresses) configured on the BIG-IP system: a VLAN for your clients, a VLAN for the WebAccelerator devices, and a VLAN for the web or application servers.

In this example, the BIG-IP LTM is the only device present on all three VLANs. This is not required, however it does provide Layer 2 segmentation of traffic.



Interior BIG-IP LTM configuration table

The following table contains a list of BIG-IP LTM configuration objects for the interior virtual server, 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.

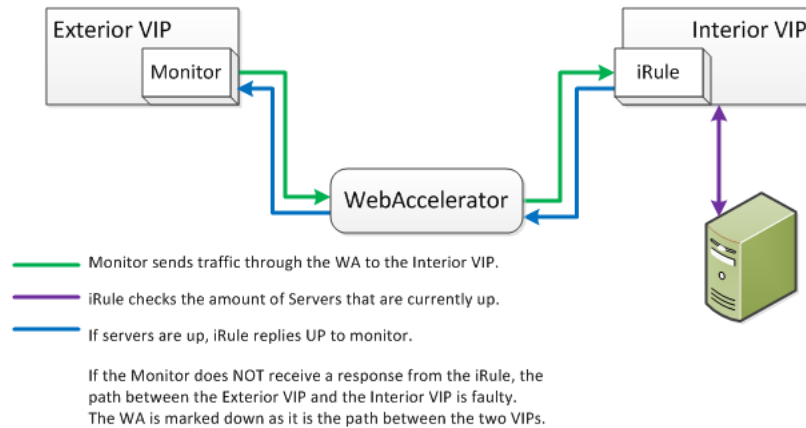
As mentioned in the introduction to this section, we are configuring the BIG-IP LTM for a generic web application in the table below. You can modify any of the BIG-IP objects (such as monitor types and profiles) for your specific application.

BIG-IP LTM Object	Non-default settings/Notes	
Health Monitor (Main tab-->Local Traffic -->Monitors)	Name	Type a unique name
	Type	Choose a monitor type specific to the application you are using. In our example, we use 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
Profiles (Main tab-->Local Traffic -->Profiles)	HTTP (Profiles-->Services)	Name: Type a unique name Parent Profile: http
	TCP LAN (Profiles-->Protocol)	Name: Type a unique name Parent Profile: tcp-lan-optimized
	Persistence (Profiles-->Persistence)	Name: Type a unique name Persistence Type: Cookie
	OneConnect (Profiles-->Other)	Name: Type a unique name Parent Profile: oneconnect
iRule (Main tab-->Local Traffic -->iRules)	See <i>Creating the iRule on page 5</i> for instructions on creating the iRule.	
Virtual Server (Main tab-->Local Traffic -->Virtual Servers)	Name	Type a unique name.
	Address	Type the IP Address for the virtual server
	Service Port	80
	Protocol Profile (client)¹	Select the LAN optimized TCP profile you created
	HTTP Profile	Select the HTTP profile you created
	OneConnect	Select the OneConnect profile you created
	SNAT Pool	Automap
	iRule	Enable the iRule you created
	Default Pool	Select the pool you created
Persistence Profile	Select the Persistence profile you created	

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

Creating the iRule

The next task is to create the iRule. This iRule is used to help monitor the health of the WebAccelerator devices. When you configure the exterior virtual server, the health monitor uses a Send String with a GET request. This iRule looks for the GET request, and if it is received marks the WebAccelerator node as up. If it is not received, the WebAccelerator node is marked down.



To create the iRule

1. On the Main tab, expand **Local Traffic**, and then click **iRules**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this profile. In our example, we type **wa-monitoring-iRule**.
4. In the **Definition** section, copy and paste the following iRule, omitting the line numbers:

Note →

The threshold of acceptable nodes down can be changed in line 3 by changing the value after '[LB::server pool]' >=' to the desired value.

```

1  when HTTP_REQUEST {
2      if { [HTTP::uri] starts_with "/monitor" } {
3          if { [active_members [LB::server pool]] >= 1 } {
4              HTTP::respond 200 content UP
5              log local0.debug "Monitor UP: [HTTP::uri]"
6          }
7      }
8      else {
9          HTTP::respond 200 content DOWN
10         log local0.debug "Monitor DOWN: [HTTP::uri]"
11     }
12 }
    
```

5. Click the **Finished** button.

This completes the interior virtual server configuration.

Configuring the WebAccelerator devices

In this section, we configure the WebAccelerator devices. In our example, the WebAccelerator devices are configured for a generic application. To get the most benefit from WebAccelerator, configure the WebAccelerator for the specific application you are using.

WebAccelerator configuration table

The following table contains a list of WebAccelerator 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.

As mentioned in the introduction to this section, we are configuring the WebAccelerator for a generic web application in the table below. You can modify any of the BIG-IP objects (such as WebAccelerator policy and HTTP class profile) for your specific application.

You must repeat this configuration for each WebAccelerator in your implementation.

BIG-IP LTM Object	Non-default settings/Notes	
Pool (Main tab-->Local Traffic -->Pools)	Name	Type a unique name
	Load Balancing Method	Round Robin
	Address	Type the IP Address of the interior BIG-IP LTM virtual server
	Service Port	80 (click Add to repeat Address and Service Port for all nodes)
Profiles (Main tab-->Local Traffic-->Profiles)	TCP LAN (Profiles-->Protocol)	Name Parent Profile tcp-lan-optimized
	OneConnect¹ (Profiles-->Other)	Name Parent Profile oneconnect
	HTTP Class (Profiles-->Protocol)	Name Parent Profile WebAccelerator httpclass Enabled
WebAccelerator Application (Main tab-->WebAccelerator-->Applications)	Application Name	Type a unique name
	Central Policy	Select the appropriate policy for your configuration. In our example, we select Level 2 Delivery .
	Requested Host	Type the Fully Qualified Domain Name (FQDN) of your application. Click Add Host to add additional hosts.
Virtual Server (Main tab-->Local Traffic -->Virtual Servers)	Name	Type a unique name.
	Address	Type the IP Address for the virtual server. This IP address needs to be within the subnet that is reachable by the LTM.
	Service Port	Type the appropriate port. In our example, we use 80 .
	Protocol Profile (client)²	Select the LAN optimized TCP profile you created
	OneConnect¹	Select the OneConnect profile you created
	SNAT Pool	Automap
	HTTP Class Profile	Enable the HTTP Class profile you created
Default Pool	Select the pool you created	

¹ Only create and apply a OneConnect profile to this virtual server if you applied a OneConnect profile on the internal LTM virtual server.

² You must select **Advanced** from the **Configuration** list for this option to appear

Repeat the configuration described in this table on each WebAccelerator in your deployment.

Configuring the BIG-IP LTM exterior virtual server

In this section, we configure the exterior virtual server on the BIG-IP LTM.

The following table contains a list of BIG-IP LTM configuration objects for the exterior virtual server, 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 Monitors (Main tab-->Local Traffic -->Monitors)	Name	Type a unique name	
	Type	Choose a monitor type specific to the application you are using. In our example, we use HTTP	
	Interval	30 (recommended)	
	Timeout	91 (recommended)	
	Send String	GET /monitor\r\n¹	
There is an additional, optional monitor that checks CPU usage of the WebAccelerator devices. See <i>Appendix: Optional EAV monitor based on CPU usage on page 8.</i>			
Pool (Main tab-->Local Traffic -->Pools)	Name	Type a unique name	
	Health Monitor	Select the monitor(s) you created above	
	Slow Ramp Time²	300	
	Load Balancing Method	Choose a load balancing method. We recommend Least Connections (Member)	
	Address	Type the IP Address of one of the WebAccelerator virtual servers you created in the previous section	
	Service Port	Type the appropriate Port. Click Add to repeat Address and Service Port for all WebAccelerator virtual servers.	
Profiles (Main tab-->Local Traffic -->Profiles)	HTTP (Profiles-->Services)	Name Parent Profile	Type a unique name http
	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
	OneConnect (Profiles-->Other)	Name Parent Profile	Type a unique name oneconnect
Virtual Server (Main tab-->Local Traffic -->Virtual Servers)	Name	Type a unique name.	
	Address	Type the IP Address for the virtual server	
	Service Port	Type the appropriate Port	
	Protocol Profile (client)²	Select the WAN optimized TCP profile you created	
	Protocol Profile (server)²	Select the LAN optimized TCP profile you created	
	HTTP Profile	Select the HTTP profile you created	
	OneConnect	Select the OneConnect profile you created	
	SNAT Pool	Automap	
Default Pool	Select the pool you created		

¹ The /monitor portion must match the URI that is being checked by the iRule. This is crucial, because without those two values matching the iRule won't work. If you have configured the iRule and monitor according to this guide, the monitor works correctly. If you modified the URI in the iRule, you must modify this Send String to match.

² You must select **Advanced** from the **Configuration** list for these options to appear

This completes the configuration.

Troubleshooting

This section contains steps to take if you are having trouble with the configuration after completing this guide.

Q: *I've configured the environment, but I can't connect to my application?*

A: Test the internal BIG-IP LTM virtual server and make sure you can reach your application. If you are unable to reach the application, check for the following on the LTM:

- Ensure the LTM is on the same VLAN as the application servers
- Ensure the LTM has a Self IP address the application servers can reach
- Verify the monitor you created for the application is properly configured

Q: *I've tested the application through the internal virtual server, but I still can not reach it through the WebAccelerator.*

A: If you are able to connect to the application using the internal virtual server, check the following on the WebAccelerator:

- Ensure the WA is on the appropriate VLAN and can be reached by the LTM
- Ensure the WA has a Self IP address that the LTM can reach
- Verify that the URL is configured correctly in the Applications section
 - » Attempt to add the IP of the WebAccelerator virtual server to the applications list and ensure you can reach the application through the WebAccelerator
- Test the full path by connecting to the external virtual server on the LTM. Make sure the application URL will let you pass to the application servers

Q: *I was able to reach the application through the WebAccelerator, but I still can't use the external virtual server.*

A: If you are unable to get through the full path but the test on the WebAccelerator was successful, check the following:

- Ensure the LTM external monitor is configured correctly
- Ensure you have a Self IP address the WA can reach on the LTM
- Ensure the SNAT Pool list is set to Automap, or you have configured a SNAT Pool and attached it to the virtual server. If you are not using SNAT, you must configure all the routing manually. See the BIG-IP documentation on manually configuring routing.

Q: *How do I turn off the monitor feedback in my logs?*

A: In the iRule, change:

```
log local0.debug "Monitor UP: [HTTP::uri]"
```


to

```
#log local0.debug "Monitor UP: [HTTP::uri]"
```

Appendix: Optional EAV monitor based on CPU usage

This appendix describes an optional monitor that can be used on the BIG-IP LTM external virtual server. This monitor is an External monitor and uses a script that provides a means to disable and enable nodes in a pool based on the CPU Usage of the specified daemon. This functionality was specifically created for monitoring the utilization of BIG-IP modules in an N+1 deployment.

In the scenario described in this guide, the traffic flow with the optional monitor looks like the following:



Prerequisites

In order to deploy this monitor successfully, you must have a few things in place before proceeding.

1. Public Key Authentication for SSH communication between BIG-IPs without passwords. For information on how to configure this SSH communications, see <http://support.f5.com/kb/en-us/solutions/public/8000/500/sol8537.html>
2. Because this monitor is being used for all WebAccelerators in this deployment, and requires a user account with administrative privileges, all of the WebAccelerators must share a user name with administrative privileges. In our example, we use **bigip**.

If your WebAccelerators do not have a administrative user account with a user name that is the same on all WebAccelerators, you must create this user on all WebAccelerator devices.

The password is not required for this script because of the SSH communication between devices described in #1 above.

3. DNS has been configured on BIG-IP system. If you have not configured the DNS settings, you can find the settings from the Main tab by expanding **System**, and then clicking **Configuration**. On the menu bar, click **Device**, and then click **DNS**.
4. Knowledge of the **pvac** Daemon. The pvac service manages HTTP and HTTPS traffic in accordance to the associated acceleration policy on the WebAccelerator.

Downloading the monitor script

First you must download and install the monitor on each BIG-IP system, create the external monitor manually that calls the script, then update the load balancing pool to use the monitor.

To download and install the monitor

1. Download the script from the following location:
<http://www.f5.com/solution-center/deployment-guides/files/pidMonitor.zip>

2. Extract the file and copy the resulting script (pidMonitor.sh) to the **/usr/bin/monitors/** directory on each of your BIG-IPs.
3. Change the permissions of the file using the following command:
chmod 755 pidMonitor.sh

The next task is to create the EAV monitor on the BIG-IP system that references the script.

To create the EAV health monitor that calls the script

Use the guidance in the following table to create a new external monitor. The table contains all of the non-default settings required for this monitor. For more information on external monitors, or for instructions on configuring the monitor, see the online help or the product documentation.

To start the monitor creation, from the BIG-IP Configuration utility Main tab, expand **Local Traffic**, click **Monitors**, and then click the **Create** button.

Monitor Field	Description/Notes	
Name	User choice.	
Type	External (the <i>Import Settings</i> field automatically selects External as well)	
Interval	User choice, but we recommend 60 .	
Timeout	User choice, but we recommend 181 .	
External Program	/usr/bin/monitors/pidMonitor.sh	
Variables	<i>Name</i>	<i>Value</i>
	Name	File name of the script. This is pidMonitor.sh unless you have changed the file name.
	User	This is a user name with admin access to the all WebAccelerator devices that will be monitored. In our example, we use bigip .
	Module	pvac (this is the daemon for WebAccelerator). You could specify a different daemon here, but for this configuration we recommend pvac.
	Limit	The CPU threshold you want to set. In our example, we use 100

This completes the monitor configuration.

Document Revision History

Version	Description
1.0	New document

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

F5 Networks,
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

