



Important: This guide has been archived. While the content in this guide is still valid for the products and versions listed in the document, it is no longer being updated and may refer to F5 or third party products or versions that have reached end-of-life or end-of-support.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>.

Deploying the BIG-IP LTM with Multiple BIG-IP WebAccelerator and ASM 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 deploying the F5 BIG-IP® Local Traffic Manager™ (LTM) with multiple BIG-IP WebAccelerator and Application Security Manager (ASM) devices. This guide shows you how to configure the BIG-IP LTM together with multiple WebAccelerator and ASM devices for fast, secure and reliable access to your applications.

This document is written for organizations deploying high volume applications that require reliable and secure access. This is accomplished by using WebAccelerator to offload from the servers using intelligent caching and compression at the first layer. The traffic is then passed onto a layer of multiple Application Security Managers to ensure security and high availability on a large scale.

The BIG-IP uses sophisticated health monitors not only to ensure that traffic is directed to available devices, but also to provide intelligent traffic management based on the utilization of the BIG-IP devices, resulting in the best possible user experience.

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, WebAccelerator and ASM	10.2.1, 10.2.2 (applies to versions 10.x)

Important! This deployment guide has been archived and will no longer be updated except for fixes. There is a newer guide for v11.4 and later you can find at <http://www.f5.com/pdf/deployment-guides/ltm-asm-aam-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) and two active ASM devices.

- 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 ASM virtual server on the BIG-IP LTM.

The BIG-IP LTM then directs the request to an available BIG-IP ASM. As the traffic passes through the ASM, it is analyzed to protect against serious security threats such as denial of service (DoS) and SQL injection, which target vulnerabilities in applications. Once the traffic has been analyzed and secured, the ASM sends the request to the internal virtual server that contains a fully customized set of monitors, policies and profiles unique to your application.

You can host the the virtual servers on the same BIG-IP LTM, or you may have separate internal and external BIG-IP LTM devices. In the following logical configuration example, we show three separate BIG-IP LTM systems for clarity. A traffic flow diagram is on the following page.

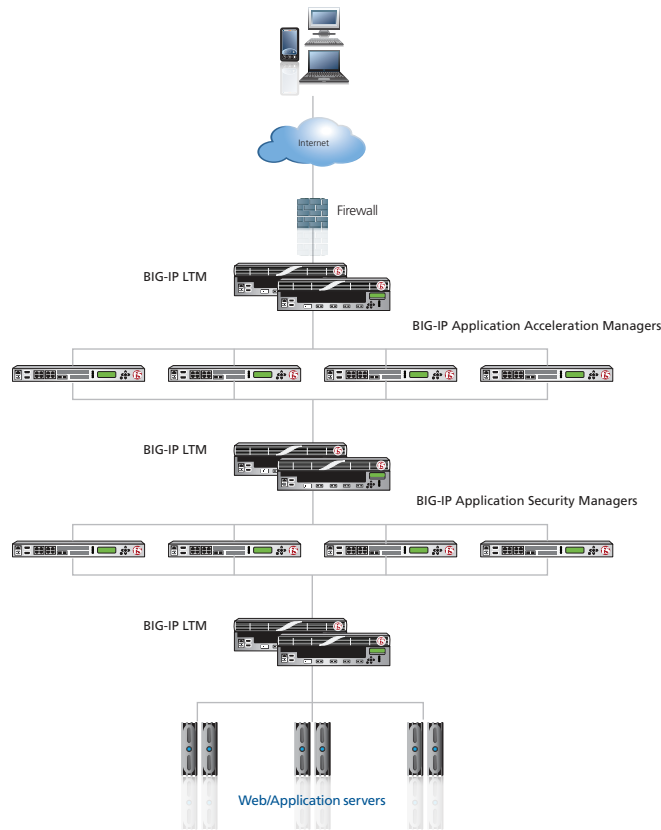


Figure 1: Logical configuration example

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

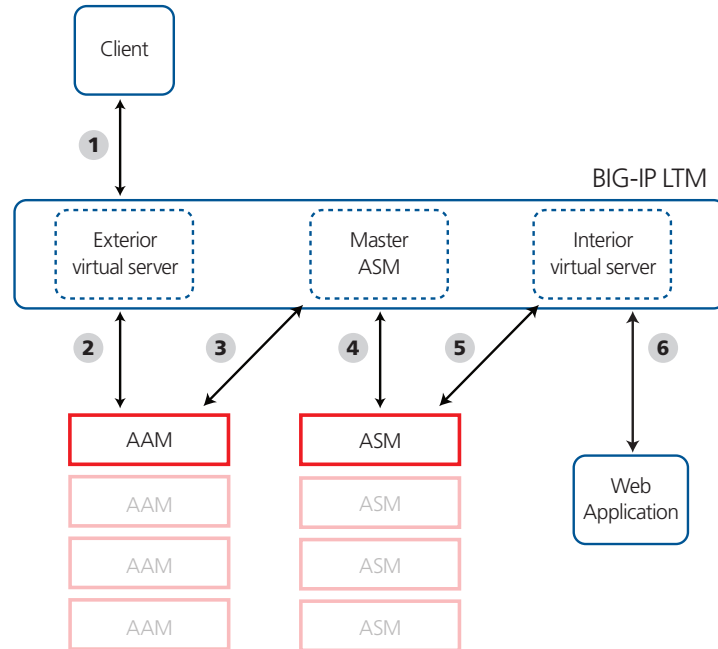


Figure 2: Configuration example

Traffic Flow

1. The client sends a request to the Web application
2. The exterior virtual server on the BIG-IP LTM receives the request, and then directs the request to an available WebAccelerator device for optimization.
3. The WebAccelerator sends the request on to the ASM virtual server on the BIG-IP LTM.
4. The BIG-IP ASM, using the application security policy, analyzes the traffic to protect the application.
5. The BIG-IP ASM device sends the request to the interior virtual server on the BIG-IP LTM.
6. The BIG-IP LTM directs the request to the appropriate application or web server depending on load balancing method and health monitoring.

Configuring the BIG-IP LTM for the internal application

In this section, we configure the virtual server for the application 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.

BIG-IP LTM configuration table for the internal application

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)	Service Port	80 (click Add to repeat Address and Service Port for all nodes)	
	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 monitoring 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 monitoring iRule

The next task is to create the iRule. This iRule is used to help monitor the health of the WebAccelerator and ASM devices. When you configure the exterior virtual server on the BIG-IP LTM, the health monitor uses a Send String with a GET request. This iRule (which you will use on both the *ASM* and *interior* virtual servers) looks for the GET request. If the iRule finds at least one node that is available (number of nodes is customizable, see *Note* below), it marks the path as UP and traffic continues to flow. If the iRule does not find a node that is available, the path is marked down, and the connection is terminated.

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 **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 BIG-IP Application Security Manager devices

In this section, we configure the BIG-IP ASM devices. In our example, the ASM devices are configured to protect a generic application. To get the most from this deployment, configure the ASM devices for the specific application you are using.

BIG-IP ASM configuration table

The following table contains a list of ASM 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 ASM for a generic web application in the table below. You can modify any of the BIG-IP objects for your specific application.

You must repeat this configuration for each ASM 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 BIG-IP LTM virtual server you created for your application in the table above.
	Service Port	80
Profiles (Main tab-->Local Traffic-->Profiles)	TCP LAN (Profiles-->Protocol)	Name Parent Profile tcp-lan-optimized
	OneConnect¹ (Profiles-->Other)	Name Parent Profile oneconnect
	HTTP (Profiles-->Services)	Name Parent Profile http
	HTTP Class (Profiles-->Protocol)	Name Parent Profile Application Security httpclass Enabled
ASM Security Policy (Main tab-->Application Security-->Web Applications)	Web Applications list	From the Web Application table, find the HTTP class you created above, and then in the Active Security Policy column, click Configure Security Policy .
	Security Policy Deployment Wizard	Follow the Security Policy wizard with information appropriate for your configuration.
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 ASM in your deployment

Configuring the ASM virtual server on the BIG-IP LTM

The next task is to create a virtual server and associated objects on the BIG-IP LTM for the ASM devices. The load balancing pool for this virtual server contains each of the ASM virtual servers you created in the preceding section.

This section covers the following two scenarios:

➤ **Fail-open**

In fail-open mode, in the unlikely event all of the ASM devices are unavailable, the BIG-IP sends the traffic directly to the application servers. This enables the ability to deploy this configuration in a production environment with zero downtime by slowly diverting traffic to the ASM devices.

While this is less secure (because the ASM devices are no longer inspecting the traffic), there is no downtime, as the requests are sent directly to the servers.

➤ **Fail-closed**

In fail-closed mode, in the unlikely event all of the ASM devices are unavailable, the request cannot complete, and eventually times out.

While this method is more secure, (because all traffic must go through the ASM devices), if none of the ASM devices are available, the end users are not granted access to the applications.

If you choose fail-closed, we recommend creating an iRule that would send the user a custom error page, and not just a 404 error. This iRule is outside the scope of this document. See [devcentral.f5.com](http://devcentral.f5.com/wiki/default.aspx/iRules/HTTP__respond.html) for more information on iRules (for example, http://devcentral.f5.com/wiki/default.aspx/iRules/HTTP__respond.html shows a possible custom error page).

The following table contains a list of BIG-IP LTM configuration objects for the ASM 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.

Note that while SNATs typically simplify the configuration, SNAT can interfere with the DoS attack prevention on BIG-IP ASM. Therefore, we do not recommend configuring a SNAT for this virtual server. The ASMs must be able to route back to the clients via the BIG-IP or have auto last hop enabled (enabled by default).

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)

This table continues on the following page

BIG-IP LTM Object	Non-default settings/Notes																	
Pool <i>(Main tab-->Local Traffic -->Pools)</i>	Name Health Monitor Load Balancing Method Priority Group Activation Address Service Port Priority	Type a unique name Select the monitor you created above Dynamic Ratio <i>For Fail-open mode only:</i> Select Less than from the list, and then in the Available Member box, type 1 . Type the IP Address of an ASM virtual server you created in the previous section. 80 <i>For Fail-open mode only:</i> In the Priority box, type 10 . <hr/> Repeat Address, Port and Priority (if applicable) for all ASM virtual servers. <i>For Fail-open mode only:</i> Repeat Address, Port and Priority to add each of the Application servers to the pool. For the Application servers only: In the Priority box, type 5 . You must give the Application servers a lower priority than the ASM virtual servers.																
Profiles <i>(Main tab-->Local Traffic -->Profiles)</i>	HTTP <i>(Profiles-->Services)</i> TCP LAN <i>(Profiles-->Protocol)</i> Persistence <i>(Profiles-->Persistence)</i> OneConnect <i>(Profiles-->Other)</i>	<table border="0"> <tr> <td>Name</td> <td>Type a unique name</td> </tr> <tr> <td>Parent Profile</td> <td>http</td> </tr> <tr> <td>Name</td> <td>Type a unique name</td> </tr> <tr> <td>Parent Profile</td> <td>tcp-lan-optimized</td> </tr> <tr> <td>Name</td> <td>Type a unique name</td> </tr> <tr> <td>Persistence Type</td> <td>Cookie</td> </tr> <tr> <td>Name</td> <td>Type a unique name</td> </tr> <tr> <td>Parent Profile</td> <td>oneconnect</td> </tr> </table>	Name	Type a unique name	Parent Profile	http	Name	Type a unique name	Parent Profile	tcp-lan-optimized	Name	Type a unique name	Persistence Type	Cookie	Name	Type a unique name	Parent Profile	oneconnect
Name	Type a unique name																	
Parent Profile	http																	
Name	Type a unique name																	
Parent Profile	tcp-lan-optimized																	
Name	Type a unique name																	
Persistence Type	Cookie																	
Name	Type a unique name																	
Parent Profile	oneconnect																	
iRule <i>(Main tab-->Local Traffic -->iRules)</i>	If you are using separate BIG-IP LTM devices for each layer, and do not have the monitoring iRule created, See <i>Creating the monitoring iRule on page 5</i> for instructions. If you are using one BIG-IP LTM device, with multiple virtual servers, there is no need to recreate the iRule.																	
Virtual Server <i>(Main tab-->Local Traffic -->Virtual Servers)</i>	Name Address Service Port Protocol Profile (client) ¹ HTTP Profile OneConnect iRule Default Pool Persistence Profile	Type a unique name. Type the IP Address for the virtual server 80 Select the LAN optimized TCP profile you created Select the HTTP profile you created Select the OneConnect profile you created Enable the iRule you created Select the pool you created Select the Persistence profile you created																

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 BIG-IP LTM ASM virtual server you created in the previous section.	
	Service Port	80	
Profiles (Main tab-->Local Traffic-->Profiles)	TCP LAN (Profiles-->Protocol)	Name Parent Profile	Type a unique name tcp-lan-optimized
	OneConnect¹ (Profiles-->Other)	Name Parent Profile	Type a unique name oneconnect
	HTTP Class (Profiles-->Protocol)	Name Parent Profile WebAccelerator	Type a unique name 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 <i>(Main tab-->Local Traffic -->Monitors)</i>	Name Type Interval Timeout Send String	Type a unique name Choose a monitor type specific to the application you are using. In our example, we use HTTP 30 (recommended) 91 (recommended) GET /monitor\r\n^1
	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 12.</i>	
Pool <i>(Main tab-->Local Traffic -->Pools)</i>	Name Health Monitor Slow Ramp Time² Load Balancing Method Address Service Port	Type a unique name Select the monitor(s) you created above 300 Choose a load balancing method. We recommend Least Connections (Member) Type the IP Address of one of the WebAccelerator virtual servers you created in the previous section Type the appropriate Port. Click Add to repeat Address and Service Port for all WebAccelerator virtual servers.
Profiles <i>(Main tab-->Local Traffic -->Profiles)</i>	HTTP <i>(Profiles-->Services)</i>	Name: Type a unique name Parent Profile: http
	TCP WAN <i>(Profiles-->Protocol)</i>	Name: Type a unique name Parent Profile: tcp-wan-optimized
	TCP LAN <i>(Profiles-->Protocol)</i>	Name: Type a unique name Parent Profile: tcp-lan-optimized
	OneConnect <i>(Profiles-->Other)</i>	Name: Type a unique name Parent Profile: oneconnect
Virtual Server <i>(Main tab-->Local Traffic -->Virtual Servers)</i>	Name Address Service Port Protocol Profile (client)² Protocol Profile (server)² HTTP Profile OneConnect SNAT Pool Default Pool	Type a unique name. Type the IP Address for the virtual server Type the appropriate Port Select the WAN optimized TCP profile you created Select the LAN optimized TCP profile you created Select the HTTP profile you created Select the OneConnect profile you created Automap 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]"
```

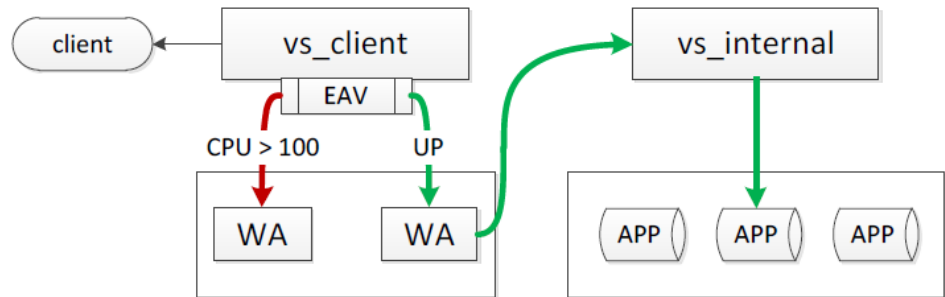
to

```
#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	Date
1.0	New guide	11-03-2011
1.1	Added a note on page 1 stating this guide has been archived and included a link to the new version.	02-04-2015

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

