Deployment Guide

# Deploying the BIG-IP LTM with Multiple BIG-IP AAM and ASM Devices

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

This document is written for organizations deploying high-volume applications based on web technologies. Local Traffic Manager (LTM) is used to direct traffic through layers of AAM and ASM devices. The AAM layer uses intelligent caching and compression to improve the application user experience while reducing the volume of requests which ASM devices and application servers actually process, minimizing application latency. The ASM layer guards application servers against malicious traffic, and the LTM provides scalability and high availability.

The BIG-IP system uses sophisticated load-balancing algorithms to provide intelligent traffic management based on the availability and performance of all devices and servers, 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

| Product | Version |
|---|---|
| BIG-IP LTM, AAM and ASM | 11.4. - 11.6 |
| Deployment guide version | 1.1 (see *Document Revision History on page 14*) |

**Important:** *Make sure you are using the most recent version of this deployment guide, available 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.*

# Contents

## Prerequisites and configuration notes

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

- You must be running BIG-IP version 11.4 or later. The configuration guidance in this document does not apply to earlier TMOS versions (a historical version of this guide may be available for use with earlier versions of BIG-IP TMOS). For versions 10.2.x and later in the 10.x branch, see *http://www.f5.com/pdf/deployment-guides/big-ip-ltm-wa-asm-dg.pdf*.
- For the configuration in this guide, you should have at least two active BIG-IP AAM devices (and not just an active/standby high availability pair) and two active BIG-IP ASM devices.
- The BIG-IP devices must be initially configured with the appropriate VLANs and Self IP addresses. For information on configuring VLANs and Self IP addresses, see the Help tab or the BIG-IP documentation.

## General Structure of the System

Like typical application servers, AAM and ASM devices are examples of resources which scale horizontally. As application traffic increases, we can increase system capacity and throughput by adding AAM and ASM devices in parallel. As we scale out AAM and ASM devices, we need to balance the load across them just as with application servers. We use LTM to do that.

We must manage connections to AAM and ASM devices in pools just as we do with application servers. A pool of similarly-configured AAM or ASM devices supports service availability even when a subset of those devices are offline. However, for performance reasons we persist each particular client's traffic to the same device when possible—to take advantage of cache locality, to reduce log-correlation effort, and to facilitate detection of subtle denial-of-service attacks. . We use LTM monitors and persistence profiles with AAM and ASM devices as with application servers.

Together with application servers, AAM and ASM become components of a system directed by LTM. As the following diagram shows, traffic flows through the system are gathered into LTM to be recognized and managed, distributed to resources such as AAM for processing, gathered again to LTM, fanned-out again, and finally load balanced to application servers.

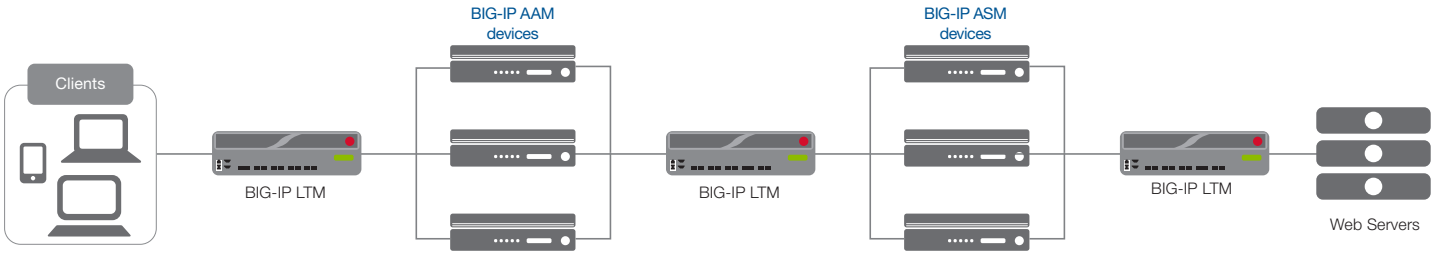


Figure 1: Logical configuration example

## Configuration example

In the configuration described in this guide, a client accesses a web application by connecting to the exterior virtual server on a BIG-IP LTM. The exterior virtual server receives each request and intelligently proxies it to an available BIG-IP AAM device in a pool of AAM devices. Each AAM device presents one AAM virtual server. The AAM virtual server uses an acceleration policy to optimize the transaction, and then sends the request to the master ASM virtual server on a BIG-IP LTM.

The master ASM virtual server then delivers the request to an available BIG-IP ASM device in a pool of ASM devices. Each ASM device presents one ASM virtual server. As the request passes through ASM it is analyzed to recognize security threats such as denial of service (DoS) and SQL injection attacks. Attacks are blocked and suspicious traffic logged for review. After analyzing and securing the request, the ASM virtual server sends it to the interior virtual server on a BIG-IP LTM. The interior virtual server distributes requests to application web servers.

You may host all the LTM virtual servers on the same device or you may use separate internal and external LTM devices. In Figure 1, we show three separate BIG-IP LTM systems for clarity. A traffic flow diagram is on the following page.

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

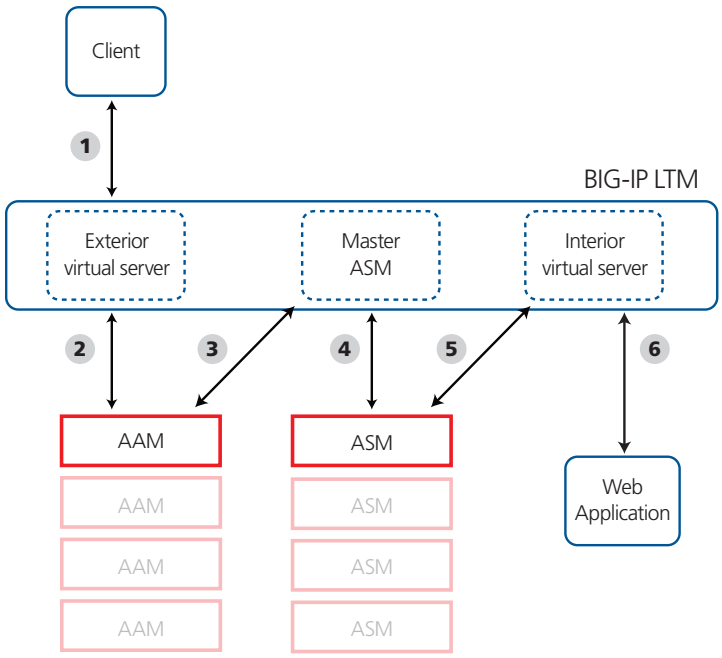


Figure 2: Configuration example

### Traffic Flow

1. The client sends a request to the web application, and the application host name resolves to the IP address of the exterior virtual server.
2. The exterior virtual server on the BIG-IP LTM receives the request, and proxies it to the virtual server of an available BIG-IP AAM device for optimization.
3. The AAM device sends the request to the master ASM virtual server on the BIG-IP LTM, if it cannot be served from cache.
4. The master ASM virtual server passes the request to the virtual server of an available BIG-IP ASM device.
5. The ASM device applies the application security policy to protect the application, and then sends the request to the interior virtual server on the BIG-IP LTM.
6. The interior virtual server directs the request to the appropriate application web server depending on load balancing method and health monitoring.

**Note**

*In this guide, the configuration begins with the internal BIG-IP configuration for the application web servers and works outward through the various layers of LTM, ASM, and AAM configuration. This way, each object has been created before it is referenced elsewhere.*

## Configuring the BIG-IP LTM for the internal application

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 load balances and shapes traffic to your web application servers. In the following procedures we use a generic HTTP web application as an example.

### Creating the monitor-response iRule

The first task is to create the iRule we use to enable upstream LTM devices to monitor the availability of a virtual server. Since BIG-IP virtual servers are designed to be (nearly) transparent to network traffic it is a bit challenging to monitor LTM, AAM, or ASM device availability with our usual "in-band" methods. We use this iRule to respond to status probes from upstream monitors.

The monitor-response iRule defined here responds to HTTP requests of the form "GET /monitor". So long as the pool for the virtual server has at least one node (server) available to handle requests, this iRule tells the outer-layer monitor that the service is UP and traffic continues to flow. When no node (server) is available, this iRule tells the calling monitor the service is DOWN, so requests will be stalled until the situation is corrected.

**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 **ir-monitor-nodecount**.
4. In the **Definition** section, copy and paste the following iRule, omitting the line numbers:

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

5. Click the **Finished** button.

The URI path **/monitor** is arbitrary but must match in the iRule and any monitor which queries it. If the real application uses **/monitor** for something, you may change the URI path in both this monitor-response iRule and any corresponding health-monitor Send String to something else which does not collide with the application, so long as you make the identical change in both places.

### BIG-IP LTM configuration table for the interior virtual server

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) to suit your specific application. You may also want to add Transport Layer Security (TLS/SSL) using BIG-IP Client and Server SSL profiles. See the BIG-IP documentation for specific details.

| BIG-IP LTM Object | Non-default settings/Notes | | |
|---|---|---|---|
| **Health Monitor** (*Main tab-->Local Traffic -->Monitors*) | ***Name*** | Type a unique name | |
| | ***Type*** | **HTTP** | |
| | ***Send String*** | **`GET / HTTP/1.1\r\nHost: webhost\r\n\r\n`**<br>**Note**: You may replace the **/** after GET with the URI of a resource in your web application which responds quickly when queried without changing any application data. You may replace **webhost** with a host name that all of your application servers recognize. You can also add a specific Receive String with the response the system should expect as a result of the Send String. | |
| **Pool** (*Main tab-->Local Traffic -->Pools*) | ***Name*** | Type a unique name | |
| | ***Health Monitor*** | Select the monitor you created above | |
| | ***Slow Ramp Time***[1] | **300** | |
| | ***Load Balancing Method*** | Choose a load balancing method. We recommend **Predictive (Member)** | |
| | ***Address*** | Type the IP Address of a node. You can optionally add a name for the node. | |
| | ***Service Port*** | **80** (click **Add** to repeat Address and Service Port for all 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*) | Be sure you have created the iRule described in *Creating the monitor-response iRule on page 5* on the device you are configuring. | | |
| **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 | |
| | ***Source Address Translation***[2] | **Auto Map**[2] | |
| | ***iRule*** | Enable the service-monitor iRule you created (**ir-monitor-nodecount**) | |
| | ***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

[2] If expecting more than 64,000 simultaneous connections per server, you must configure a SNAT Pool. See the BIG-IP documentation on configuring SNAT Pools.

## Configuring the BIG-IP Application Security Manager devices

In this section, we configure the BIG-IP ASM devices. Each ASM device supports one virtual server with an Application Security Policy enabled on it. In our example, we configure ASM to protect a generic application. To get the most from this deployment, tune your Application Security Policy to your specific application. See the BIG-IP ASM documentation for specific details.

### 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. 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 | | |
|---|---|---|---|
| **Health Monitor** (*Main tab-->Local Traffic -->Monitors*) | ***Name*** | Type a unique name | |
| | ***Type*** | **HTTP** | |
| | ***Send String*** | **`GET /monitor HTTP/1.1\r\nHost: webhost\r\n\r\n`** | |
| | ***Receive String*** | **UP** | |
| | ***Receive Disable String*** | **DOWN** | |
| **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 you created in the previous section. | |
| | ***Service Port*** | **80** | |
| **Profiles** (*Main tab--> Local Traffic-->Profiles*) | ***TCP LAN*** (*Profiles-->Protocol*) | Name | Type a unique name |
| | | Parent Profile | **tcp-lan-optimized** |
| | ***OneConnect***[1] (*Profiles-->Other*) | Name | Type a unique name |
| | | Parent Profile | **oneconnect** |
| | ***HTTP*** (*Profiles-->Services*) | Name | Type a unique name |
| | | Parent Profile | **http** |
| | | Accept XXF | **Enabled** |
| **Logging Profiles** (*Main tab-->Security--> Event Logs--> Logging Profiles*) | ***Name*** | Type a unique name | |
| | ***Application Security*** | **Enabled** (checked). The Application Security details appear. | |
| | ***Remote Storage*** | If you have a remote logging or SIEM server, enable Remote Storage and enter appropriate configuration data. | |
| | You can alternatively create a logging profile using an iApp template, see *https://devcentral.f5.com/wiki/iApp.F5-Remote-Logging-iApp.ashx* | | |
| **iRule** (*Main tab-->Local Traffic -->iRules*) | Be sure you have created the iRule described in *Creating the monitor-response iRule on page 5* on the device you are configuring. | | |
| **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 | |
| | ***Source Address Translation***[2] | **Auto Map**[2] | |
| | ***iRule*** | Enable the service-monitor iRule you created (**ir-monitor-nodecount**) | |
| | ***Default Pool*** | Select the pool you created | |
| | ***Security Log Profile*** | Click **Enabled**, and then select the logging profile you created. | |
| *Adding the Logging profile to the virtual server* | After you have created the virtual server, from the *Local Traffic-->Virtual Servers* list, click the name of the virtual server you just created. On the menu bar, click **Security**. In the Policy Settings area, from the **Log Profile** list, select **Enabled**, and then move the Logging Profile you created to the **Selected** box. Click **Update**. | | |

| BIG-IP LTM Object | Non-default settings/Notes | |
|---|---|---|
| **ASM Security Policy** (*Main tab-->Security--> Application Security--> Security Policies*) | **Active Security Policies** | Click **Create**. The Deployment Wizard opens.<br>For the **Local Traffic Deployment Scenario**, select **Existing Virtual Server**. Continue through the wizard to create and attach an ASM Security Policy to the virtual server you just created. |
| | **Security Policy Deployment Wizard** | To achieve a rapid initial deployment, on the **Deployment Scenario** page, select **Create a security policy manually or use templates**, and then for **Application Ready Security Policy** choose **Rapid Deployment Security Policy**. When you Configure Attack Signatures, you may be able to add extra Systems. Before you Save the results of the wizard, consider whether to choose the Enforcement Mode "Transparent" or "Blocking": Transparent will log but not block attacks.<br>***NOTE:*** After creating a security policy on one ASM device, you can export it and then import it to additional ASM devices to ensure consistency and save effort. |

[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

Remember to repeat the configuration described in this table on each ASM in your deployment.

## Configuring the Master ASM virtual server on the BIG-IP LTM

Use the following guidance to create a master ASM virtual server and associated objects on the BIG-IP LTM. This virtual server will load balance requests to the ASM virtual servers (one per ASM device) you created in the preceding section.

This section covers the following two scenarios:

- **Fail-unsecured**
  In fail-unsecured mode, when no ASM devices are available the LTM sends traffic directly to the web application servers. This lets you deploy this configuration in a production environment with zero downtime by bringing ASM devices on or offline gradually.

  This mode is less secure because traffic may sometimes reach the application web servers without being secured by ASM, but it avoids downtime when connectivity problems or administrative work affect ASM availability.

- **Fail-secure**
  In fail-secure mode, when no ASM devices are available requests will not be processed.

  While this method is more secure because all traffic must go through the ASM devices, if no ASM devices are available then end users will not be able to use the web application.

  If you choose fail-secure you may wish to add an iRule to the exterior virtual server to send application users a "friendly" error message when the application is unavailable (otherwise they simply receive no response). A minimal example of such an iRule is:

```
1  when HTTP_REQUEST {
2     if { [active_members [LB::server pool]] < 1 } {
3        HTTP::respond 503 \
4            content "Application [HTTP::host] temporarily unavailable. \
5            Please try again later."  Content-Type "text/plain; charset=UTF-8"
6        return
7     }
8  }
```

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.

| BIG-IP LTM Object | Non-default settings/Notes | |
|---|---|---|
| **Health Monitor** (*Main tab-->Local Traffic -->Monitors*) | ***Name*** | Type a unique name |
| | ***Type*** | **HTTP** |
| | ***Send String*** | **`GET /monitor HTTP/1.1\r\nHost: webhost\r\n\r\n`** |
| | ***Receive String*** | **`UP`** |
| | ***Receive Disable String*** | **`DOWN`** |
| **Pool** (*Main tab-->Local Traffic -->Pools*) | ***Name*** | Type a unique name |
| | ***Health Monitor*** | Select the monitor you created above |
| | ***Load Balancing Method*** | **Observed (Member)** |
| | ***Priority Group Activation*** | *For **Fail-open** mode only:* Select **Less than** from the list, and then in the **Available Member** box, type **1**. |
| | ***Address*** | Type the IP Address of an ASM virtual server you created in the previous section. |
| | ***Service Port*** | **80** |
| | ***Priority*** | *For **Fail-open** mode only:* In the **Priority** box, type **10**. |
| | Repeat Address, Port and Priority (if applicable) for all ASM virtual servers (devices).<br><br>*For **Fail-open** mode only:* Use Address, Port, and Priority to add each actual application server to the pool. For application servers (only) in the Priority box enter 5. *You must give every application server a lower priority than any ASM virtual server.* | |
| **Profiles** (*Main tab-->Local Traffic -->Profiles*) | ***HTTP*** (*Profiles-->Services*) | Name: Type a unique name<br>Parent Profile: **http**<br>Accept XFF: **Enable** |
| | ***TCP LAN*** (*Profiles-->Protocol*) | Name: Type a unique name<br>Parent Profile: **tcp-lan-optimized** |
| | ***Persistence*** (*Profiles-->Persistence*) | Name: Type a unique name<br>Persistence Type: **Cookie** |
| | ***OneConnect*** (*Profiles-->Other*) | Name: Type a unique name<br>Parent Profile: **oneconnect** |
| **iRule** (*Main tab-->Local Traffic -->iRules*) | If you are using separate BIG-IP LTM devices for each layer, and do not have the monitor-response iRule created, see *Creating the monitor-response iRule on page 5* for instructions on creating this iRule.<br><br>If you are using one BIG-IP LTM device with multiple virtual servers, there is no need to recreate 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 |
| | ***iRule*** | Enable the monitoring iRule you created (ir-monitor-nodecount) |
| | ***Default Pool*** | Select the pool you created |
| | ***Persistence Profile*** | Select the Persistence profile you created |

For specific instructions on configuring individual objects, see the online help or product manuals.

## Configuring the BIG-IP AAM devices

In this section we configure the AAM devices. In our example, AAM is configured for a generic application. To get the most benefit from AAM, configure AAM features for your specific application.

### BIG-IP AAM configuration table

The following table contains a list of AAM 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 AAM for a generic web application in the table below. You can modify any of the BIG-IP objects (such as AAM policy and HTTP Compression profile) for your specific application.

You must repeat this configuration for each BIG-IP AAM in your implementation.

| BIG-IP LTM Object | Non-default settings/Notes | | |
|---|---|---|---|
| **Web Application** (*Main tab-->Acceleration--> Web Applications-->Applications*) | **Application Name** | Type a unique name | |
| | **Policy** | **Generic Policy - Enhanced** | |
| | **Requested Host** | ***** (asterisk)<br>Use the **Add Host** button to included more host names. | |
| **Accelerator Profile** (*Main tab-->Acceleration--> Profiles--> Web Acceleration*) | **Name** | Type a unique name | |
| | **Parent Profile** | Must be: **optimized-acceleration** | |
| | **AM Applications** | Select the Web Application you just defined and then click Enable | |
| **Health Monitor** (*Main tab-->Local Traffic -->Monitors*) | ***Name*** | Type a unique name | |
| | ***Type*** | **HTTP** | |
| | ***Send String*** | **`GET /monitor HTTP/1.1\r\nHost: webhost\r\n\r\n`** | |
| | ***Receive String*** | **UP** | |
| | ***Receive Disable String*** | **DOWN** | |
| **Pool** (*Main tab-->Local Traffic -->Pools*) | ***Name*** | Type a unique name | |
| | ***Health Monitor*** | Select the monitor you created above | |
| | ***Load Balancing Method*** | **Round Robin** | |
| | ***Address*** | Type the IP Address of the ***Master ASM virtual server*** you created in the previous section | |
| | ***Service Port*** | **80** | |
| **Profiles** (*Main tab--> Local Traffic-->Profiles*) | ***TCP LAN*** (*Profiles-->Protocol*) | Name | Type a unique name |
| | | Parent Profile | **tcp-lan-optimized** |
| | ***OneConnect***[1] (*Profiles-->Other*) | Name | Type a unique name |
| | | Parent Profile | **oneconnect** |
| | ***HTTP*** (*Profiles-->Protocol*) | Name | Type a unique name |
| | | Parent Profile | **httpclass** |
| | | Accept XFF | **Enabled** |
| | ***HTTP Compression*** (*Profiles-->Protocol*) | Name | Type a unique name |
| | | Parent Profile | **wan-optimized-compression** |
| **iRule** (*Main tab-->Local Traffic-->iRules*) | See *Creating the monitor-response iRule on page 5* for instructions. | | |

| BIG-IP LTM Object | Non-default settings/Notes | |
|---|---|---|
| **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; this is typically port **80** |
| | ***Protocol Profile (client)***[2] | Select the LAN optimized TCP profile you created |
| | ***HTTP Profile*** | Select the HTTP profile you created |
| | ***OneConnect***[1] | Select the OneConnect profile you created |
| | ***Secure Address Translation*** | **Auto Map** |
| | ***HTTP Compression Profile*** | Enable the HTTP Compression profile you created |
| | ***Web Acceleration Profile*** | Enable the Web Acceleration 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 AAM device 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*** | **HTTP** | |
| | ***Send String*** | **`GET /monitor HTTP/1.1\r\nHost: webhost\r\n\r\n`** | |
| | ***Receive String*** | **UP** | |
| | ***Receive Disable String*** | **DOWN** | |
| **Pool** (*Main tab-->Local Traffic -->Pools*) | ***Name*** | Type a unique name | |
| | ***Health Monitor*** | Select the monitor(s) you created above | |
| | ***Slow Ramp Time***[2] | **300** | |
| | ***Load Balancing Method*** | Choose a load balancing method. We recommend **Predictive (Member)** | |
| | ***Address*** | Type the IP Address of an AAM virtual server you created in the previous section | |
| | ***Service Port*** | Type the appropriate Port.<br>Click **Add** to repeat Address and Service Port for all AAM virtual servers. | |
| **Profiles** (*Main tab-->Local Traffic -->Profiles*) | ***HTTP*** (*Profiles-->Services*) | Name | Type a unique name |
| | | Parent Profile | **http** |
| | | Insert X-Forwarded-For | **Enabled** |
| | ***TCP WAN*** (*Profiles-->Protocol*) | Name | Type a unique name |
| | | Parent Profile | **tcp-wan-optimized** |
| | ***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** |
| **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)***[1] | Select the WAN optimized TCP profile you created | |
| | ***Protocol Profile (server)***[1] | Select the LAN optimized TCP profile you created | |
| | ***HTTP Profile*** | Select the HTTP profile you created | |
| | ***OneConnect*** | Select the OneConnect profile you created | |
| | ***Secure Address Translation*** | **Auto Map** | |
| | ***Default Pool*** | Select the pool you created | |
| | ***Default Persistence Profile*** | Select the Cookie persistence profile you created | |

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

This completes the configuration.

## Troubleshooting

This section contains advice on resolving configuration problems after completing this guide.

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

**A:** Test the interior 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 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 interior virtual server, so why am I unable to reach it through the exterior virtual server?*

**A:** If you can connect to the application using the interior virtual server, work outwards checking each ASM device (connect directly to its virtual server), the master ASM virtual server, each AAM device (connect directly to its virtual server), and finally the exterior virtual server.

- Ensure each BIG-IP has a Self IP address reachable from the next layer inward
- Ensure each virtual server is on the appropriate VLAN and is reachable from the next layer outward
- Check that each virtual server has the monitor-response iRule attached, and that the special URI matches in the monitor Send String and the iRule.
- On AAM devices, verify that the Requested Host name is configured correctly in the Web Application. You may use the asterisk wildcard as shown.

**Q:** *I was able to reach the application through the AAM, so why am I unable to use the external virtual server?*

**A:** Check the following:

- Ensure the LTM exterior monitor is configured correctly (and the monitor-response iRule on each AAM virtual server)
- Ensure the Secure Address Translation list is set to Auto Map, 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 on debugging log messages from the monitor-response iRule?*

**A:** In the iRule, change:

```
#log local0.debug "Monitor UP:..."

to

log local0.debug "Monitor UP:..."
```

## Document Revision History

| Version | Description | Date |
|---|---|---|
| 1.0 | New guide | 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