

# Deploying Layer 7 Gateways with F5 BIG-IP Local Traffic Manager



Layer 7 Technologies

## *Deployment Guide*



## Table of Contents

Introducing the F5 BIG-IP LTM and Layer 7 SOA Gateway configuration .....	3
Prerequisites and configuration notes .....	3
Configuration example .....	4
Configuring the BIG-IP LTM system for deployment with Layer 7 SOA Gateways.....	4
Configuring the Layer 7 Gateways for clustered deployment .....	5
Preparing the Layer 7 Gateway cluster .....	5
Connecting to the Gateway cluster using the Layer 7 Policy Manager .....	5
Creating a user for health checks from the BIG-IP LTM .....	5
Assigning a role to the health check user.....	6
Connecting to the BIG-IP LTM device.....	7
Connecting to the BIG-IP web-based configuration utility.....	7
Configuring the BIG-IP LTM system for Layer 7 Gateways.....	8
Creating the custom HTTPS health monitor.....	8
Creating the Pool.....	10
Creating the Virtual Server .....	13

# Introducing the F5 BIG-IP LTM and Layer 7 SOA Gateway configuration

Welcome to the F5 BIG-IP Local Traffic Manager (LTM) and Layer 7 SOA Gateway deployment guide. This guide provides step-by-step instructions on configuring BIG-IP LTM for deployment with multiple SOA Gateways for greater scalability and availability.

Layer 7 Technologies is a leading provider of API security and governance for Service Oriented Architectures, Cloud application deployments, and developer/partner-facing APIs. Through this award winning line of SecureSpan and CloudSpan family of API gateways and management products, Layer 7 is helping organizations control how they expose their data and applications to outside divisions, partners, mobile developers and cloud services.

For more information on Layer 7 Technologies, see <http://www.layer7.com/products>.

F5 Networks, Inc., the global leader in Application Delivery Networking (ADN), helps the world's largest enterprises and service providers realize the full value of virtualization, cloud computing, and on-demand IT. F5® solutions help integrate disparate technologies to provide greater control of the infrastructure, improve application delivery and data management, and give users seamless, secure, and accelerated access to applications from their corporate desktops and smart devices. An open architectural framework enables F5 customers to apply business policies at “strategic points of control” across the IT infrastructure and into the public cloud. F5 products give customers the agility they need to align IT with changing business conditions, deploy scalable solutions on demand, and manage mobile access to data and services. Enterprises, service and cloud providers, and leading online companies worldwide rely on F5 to optimize their IT investments and drive business forward.

For more information on F5 products, see <http://www.f5.com/products>.


## Prerequisites and configuration notes

The following are general prerequisites for this deployment:

- ◆ This guide was tested using the SecureSpan SOA Gateway version 6.0; initial deployment guidelines should hold true for other products in the SecureSpan product line.
- ◆ This guide was tested using BIG-IP LTM version 10.2.
- ◆ This guide was written with the assumption that you are familiar with both the BIG-IP LTM system and the Layer 7 SOA Gateway product lines. For more information on configuring these products, consult the appropriate documentation.

## Configuration example

The BIG-IP LTM system provides intelligent traffic management, failover, availability, and scalability for Layer 7 appliances. Through advanced health checking capabilities, the BIG-IP LTM system recognizes when gateway cluster nodes are unavailable and directs traffic to an available resource.



*Figure 1: F5 – Layer 7 logical configuration example*

## Configuring the BIG-IP LTM system for deployment with Layer 7 SOA Gateways

To configure the BIG-IP LTM system for directing incoming data traffic to Layer 7 Gateways, complete the following tasks:

- *Configuring Layer 7 Gateways for clustered deployment, on page 5.*
- *Configuring the BIG-IP LTM system for Layer 7 Gateways, on page 8.*

We recommend you save your existing BIG-IP LTM configuration before you begin the procedures in this Deployment Guide. For information on

backing up or restoring a BIG-IP LTM configuration, refer to the appropriate BIG-IP LTM manual, available from <http://ask.f5.com/>

## Configuring the Layer 7 Gateways for clustered deployment

To ensure that Gateway nodes are identically configured and prepared for traffic distribution by the BIG-IP LTM system, they must be configured as a Gateway cluster. Follow the instructions available in the *Layer 7 Installation & Maintenance Manual* for configuring the individual nodes and initializing database replication on each of the Gateway database nodes. Then use the following instructions to prepare the cluster for health checking and load distribution by the BIG-IP LTM.

### Preparing the Layer 7 Gateway cluster

In order to retrieve information about the health of individual Gateway nodes, the BIG-IP LTM must have access to the Gateway health check mechanism. This access can be granted for the entire cluster using the Layer 7 Policy Manager.

The following instructions assume that the Policy Manager has been installed and has network connectivity to the Gateway nodes. It also assumes that the health check will arrive via the HTTPS port on the Layer 7 Gateway, and that credentials will need to be provided. For simple HTTP health checks, or to remove the requirement for incoming credentials, modify the `pingServlet.mode` Gateway property as described in the Layer 7 Installation documentation.

#### Connecting to the Gateway cluster using the Layer 7 Policy Manager

1. Open the Layer 7 Policy Manager. The **Login** dialog will appear.
2. In the **User Name** and **Password** fields, provide credentials that have access to operate the Gateway.
3. In the **Gateway** field, type the hostname of the Gateway cluster, as defined during installation.
4. Click **OK** to establish the connection.

#### Creating a user for health checks from the BIG-IP LTM

1. In the upper left portion of the Policy Manager, click on the **Identity Providers** tab.
2. Right-click on the **Internal Identity Provider** and choose **Create User** from the menu. The **Create Internal User** dialog will appear.

3. Type **F5HealthCheck** in the **User Name** field.
4. Select a password for this user and type it in the **Password** and **Confirm** fields.
5. Uncheck the **Force Password Change** box.
6. Click the **Create** button to confirm the user creation.

The dialog should appear similar to the following screen shot.

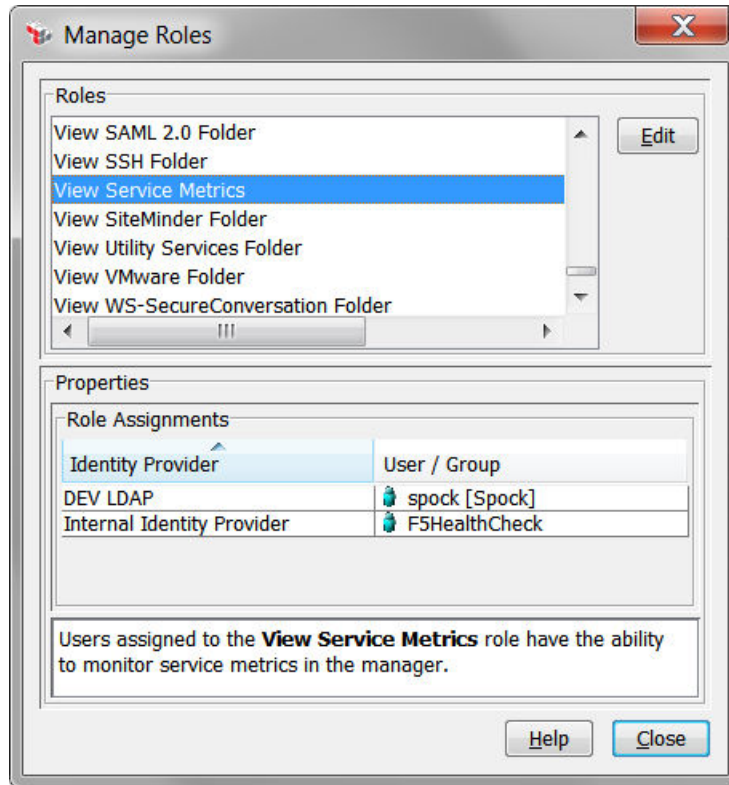
*Figure 2: Creating a new internal user*

### Assigning a role to the health check user

1. In the Policy Manager menu bar, select **Tasks** and choose **Manage Roles** from the drop-down list. The Manage Roles dialog will appear.
2. There are five predefined roles that include access to the health check service: *Administrator*, *Operator*, *Manage Cluster Status*, *View Audit Records and Logs*, and *View Service Metrics*. Select one of these from the **Roles** section of the dialog, and then click **Edit**. The **Edit RoleName Role** dialog will appear. In our example, we selected the View Service Metrics role.
3. In the **Role Assignments** section, click **Add**. The **Search Identity Provider** dialog will appear.
4. Select **Internal Identity Provider** from the **Search** field.
5. Optionally select **Type** and **Name** parameters to filter the results; the default of **ALL Equals <blank>** will return all users within the Internal Identity Provider.
6. Click **Search** to retrieve user information.
7. Select the **F5HealthCheck** user from the **Search Results**.
8. Click **Select** at the bottom of the dialog to add this user to the selected Role.

- Click **OK** to accept the changes to the Role and return to the Manage Roles dialog.

The chosen role should now have the health check user as one of its members, as shown in the following screen shot. This will allow the BIG-IP LTM to securely connect as the defined user.



*Figure 3: Adding a new user to a role*

## Connecting to the BIG-IP LTM device

Use the following procedure to access the BIG-IP LTM web-based Configuration utility using a web browser. This guide assumes that the BIG-IP LTM device has previously been installed, licensed, provisioned, and addressed on the appropriate network subnets.

### Connecting to the BIG-IP web-based configuration utility

- In a compatible browser, type the following URL: [https://<BIG-IP\\_Admin\\_IP>](https://<BIG-IP_Admin_IP>).

2. Type your username and password and click OK to access the Welcome screen.

The Welcome screen provides easy access to documentation, plugins, relevant support links, and SNMP MIBs for off-device monitoring. This interface also allows for the modification of networking details, traffic routing policies, and system-level configuration parameters.

## Configuring the BIG-IP LTM system for Layer 7 Gateways

The first task is to choose a new IP address, host name and port that will be used as the IP, name and port for your incoming connections. Register this information with your company's DNS and network administration groups before proceeding.

For example, let's assume there is a Layer 7 Gateway cluster currently with four Gateway nodes that will be used in the high availability architecture. In our examples, we use the name `Services.companyname.com`, choose port number 8443, and assign the Services hostname an IP Address of 192.168.9.100.

The BIG-IP LTM will:

- Listen for incoming connection requests directed to the BIG-IP LTM Virtual Server host name, address, and port.
- Use a BIG-IP LTM Health Monitor to determine which Gateway nodes are available to accept incoming connection requests.
- Distribute incoming connection requests across the Gateway nodes defined in a BIG-IP LTM Pool.

This requires several configuration steps on the BIG-IP LTM, which are discussed in the subsequent sections:

- Create a custom HTTPS monitor
- Create a pool of Gateway nodes
- Create a virtual server

Each Gateway node that will become a member in the Gateway cluster must possess its own hostname and IP address; these values should be noted before configuring the BIG-IP LTM to distribute incoming connections across the Gateways.

### Creating the custom HTTPS health monitor

Defining a custom health monitor will allow the BIG-IP LTM to determine if each member of the Gateway cluster is capable of processing an incoming connection request. Layer 7 Gateways provide a specific health check URL that determines availability of both the nodes and of the underlying cluster storage, thus demonstrating service availability.

1. From the BIG-IP configuration interface, expand the **Local Traffic** section of the left side menu, and then click **Monitors**.
2. Click **Create** in the upper right corner of the Monitor List. The New Monitor screen appears.
3. In the **Name** field, type a name for the monitor, such as **SSG\_HTTPS\_Monitor**.
4. From the **Type** list, select **HTTPS** to retrieve the defaults for https-style health monitors.
5. In the **Configuration** section, choose interval and timeout lengths (in seconds). In order to provide reasonable failure detection while avoiding excessive network traffic or server load, we recommend an **Interval** of **40** and a **Timeout** of **121**.
6. In the **Send String** field, type the following string:  
**GET /ssg/ping\r\n**
7. In the **Receive String** field, type the following string:  
**OK**
8. In the **User Name** and **Password** fields, type the credentials of the user making the call. To define permissions for this operation, see *Configuring the Layer 7 Gateway for clustered deployment*, on page 5. In our example, we have defined the username as **F5HealthCheck**.
9. The rest of the monitor settings are optional; configure as appropriate for your deployment.
10. Click **Finished** to complete the Health Monitor configuration.

Upon completion, the created Health Monitor should resemble the following screen shot.

The screenshot shows the configuration page for a monitor named 'SSG\_HTTPS\_Monitor'. The breadcrumb navigation is 'Local Traffic >> Monitors >> SSG\_HTTPS\_Monitor'. There are two tabs: 'Properties' (selected) and 'Instances'. The 'General Properties' section includes:

Name	SSG_HTTPS_Monitor
Partition	Common
Type	HTTPS

The 'Configuration' section is set to 'Basic'. The configuration parameters are:

Interval	40 seconds
Timeout	121 seconds
Send String	GET /sag/ping\r\n
Receive String	OK
Receive Disable String	
Cipher List	DEFAULT:+SHA:+3DES:+kEDH
User Name	F5HealthCheck
Password	••••••••
Reverse	<input type="radio"/> Yes <input checked="" type="radio"/> No
Transparent	<input type="radio"/> Yes <input checked="" type="radio"/> No

*Figure 4: Defining a new Monitor*

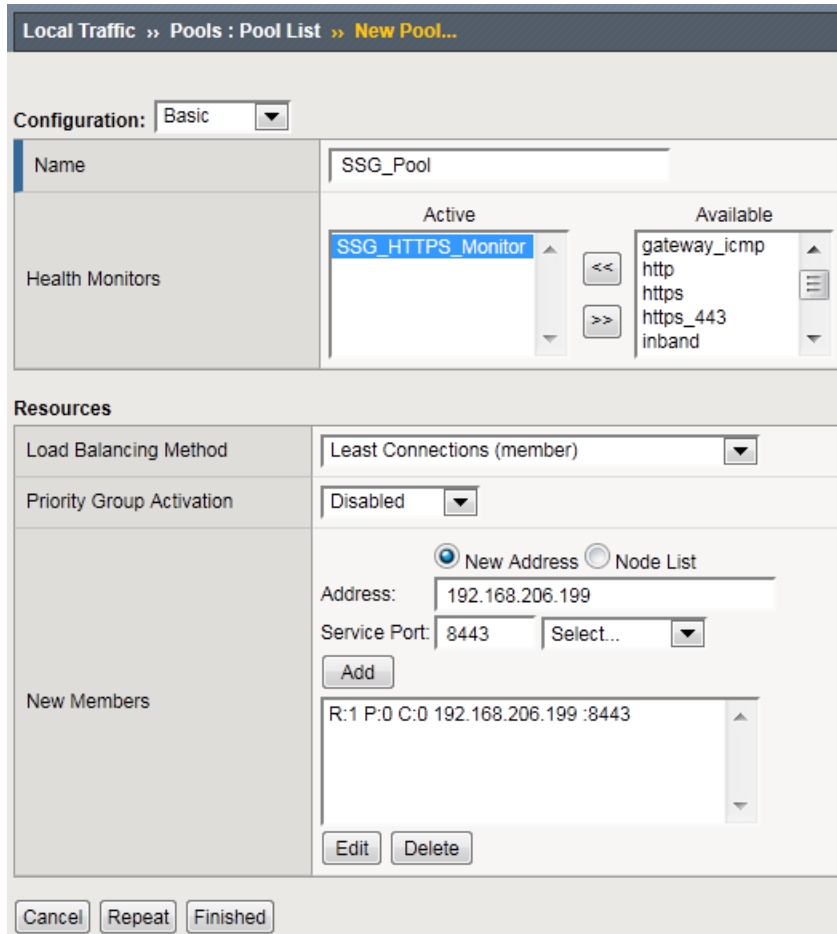
## Creating the Pool

Defining a pool of Layer 7 Gateways allows them to achieve true clustering, in that any of the included nodes can respond to a given request. Pool configuration includes a load balancing algorithm for distribution of requests, and an associated health monitor for proactive health determination.

1. From the BIG-IP configuration interface, expand the **Local Traffic** section of the left side menu, and then click **Pools**.

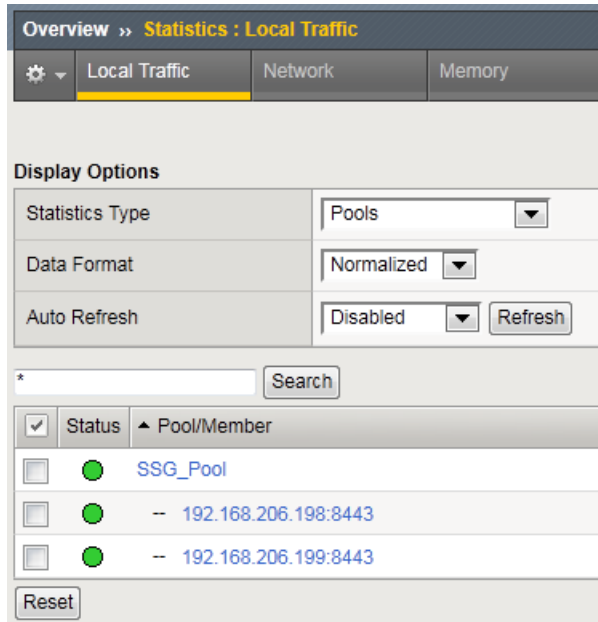
2. Click **Create** in the upper right corner of the Pool List. The New Pool screen appears.
3. In the **Name** field, type a name for the pool, such as **SSG\_Pool**.
4. In the **Health Monitors** field, select the monitor created in the *Creating the custom HTTPS health monitor* section, and click the left-pointing arrows to move it from Available to Active. In our example, we select **SSG\_HTTPS\_Monitor**.
5. In the **Load Balancing Method** section, choose your preferred load balancing algorithm. Your selection may depend on network configuration or expected traffic patterns; we select **Least Connections (member)** for our example.
6. In the **New Members** section, add the first member by typing the IP address of a Gateway node in the **Address** field. In our example, we type **192.168.206.199**.
7. In the **Service Port** field, type the port for the HTTPS listener on your Layer 7 Gateway. In our example, we typed the default SSL port of **8443**.
8. Click **Add** to add the first member to the pool.
9. Repeat steps 6-8 for each additional member of the pool. When all members have been added, click **Finished** to complete the Pool configuration.

The following screen shot shows our example pool after the first member has been added.



*Figure 5: Defining a new Pool*

10. Test the successful configuration and health check of pool members. From the BIG-IP configuration interface, expand the **Local Traffic** section of the left side menu, hover over the **Pools** option, and then click **Statistics**.
11. Confirm that the pool and all of its members are marked as healthy, which is designated by a green circle. Our example pool is shown in the following screen shot.



*Figure 6: Verifying Pool member health*

## Creating the Virtual Server

The Layer 7 Gateways that have been added to the pool define a cluster of message processing nodes. Incoming messages will traverse the BIG-IP LTM through a Virtual Server, which allows directed access between external-facing and internal-facing networks.

1. From the BIG-IP configuration interface, expand the **Local Traffic** section of the left side menu, and then click **Virtual Servers**.
2. Click **Create** in the upper right corner of the Virtual Server List. The New Virtual Server screen appears.
3. In the **Name** field, type a name for the server, such as **SSG\_HTTPS\_VS**.
4. For a single IP, select **Host** and type a specific IP address in the **Address** field.
5. In the **Service Port** field, type the port on the external network that will be forwarded to the internal Gateways. We recommend that ports used in the Virtual Server map directly to the same ports on the Layer 7 Gateway. This is mandatory if service clients will be using the XML VPN Client or any client-side tool (such as the Route via SecureSpan Bridge Assertion) that takes advantage of the Gateway's policy download capability.
6. If the network subnet used for the virtual server is the same subnet used by the Layer 7 Gateways, then change the **SNAT Pool** field to **Auto Map**. If the virtual server is defined on a different subnet, then leave the default value of **None**.

7. In the **Default Pool** field, select the Gateway Pool from the drop-down list. In our example, **SSG\_Pool** is selected.
8. Click **Finished** at the bottom of the screen to complete the Virtual Server configuration.

The upper portion of the Virtual Server page should resemble the following screen shot.

Local Traffic » Virtual Servers : Virtual Server List » New Virtual Server...

**General Properties**

Name	SSG_HTTPS_VS	
Destination	Type:	<input checked="" type="radio"/> Host <input type="radio"/> Network
	Address:	192.168.9.100
Service Port	8443	Other: <input type="text"/>
State	Enabled <input type="text"/>	

Configuration:

Type	Standard <input type="text"/>
Protocol	TCP <input type="text"/>
OneConnect Profile	None <input type="text"/>
NTLM Conn Pool	None <input type="text"/>
HTTP Profile	None <input type="text"/>
FTP Profile	None <input type="text"/>
SSL Profile (Client)	None <input type="text"/>
SSL Profile (Server)	None <input type="text"/>
Diameter Profile	None <input type="text"/>
SIP Profile	None <input type="text"/>
VLAN and Tunnel Traffic	All VLANs and Tunnels <input type="text"/>
SNAT Pool	None <input type="text"/>

*Figure 7: Defining a new virtual server*

Upon completion, the Virtual Server will be ready to accept requests destined for the Layer 7 Gateway cluster.

## About Layer 7 Technologies

With more than 150 customers across 6 continents, and successful partnerships with some of the largest ISVs and resellers in the industry, Layer 7 Technologies is the leader in SOA and cloud security and governance. Our award-winning SecureSpan™ family of XML Gateways feature sophisticated runtime governance, enterprise-scale management and industry-leading XML security. Our CloudSpan™ family enables enterprises and service providers to securely consume cloud services, as well as protect and control their own applications deployed in public and private clouds. Founded in 2002, Layer 7 has a history of helping organizations address their security, visibility and governance issues by enabling them to control, manage and adapt their Web services, no matter the deployment model – in the enterprise or in the cloud.

## Contact Layer 7 Technologies

Layer 7 Technologies welcomes your questions, comments, and general feedback.

**Support:**

[support@layer7.com](mailto:support@layer7.com)

**Email:**

[info@layer7.com](mailto:info@layer7.com)

**Web Site:**

[www.layer7.com](http://www.layer7.com)

**Phone:**

604-681-9377

1-800-681-9377 (toll free)

**Fax:**

604-681-9387

**Address:**

US Office

1200 G Street, NW, Suite 800

Washington, DC 20005

Canada Office

Suite 405-1100 Melville Street

Vancouver, BC

V6E 4A6 Canada

## Legal Information

Copyright © 2011 by Layer 7 Technologies, Inc. ([www.layer7.com](http://www.layer7.com)). Contents confidential. All rights reserved. SecureSpan™ is a registered trademark of Layer 7 Technologies, Inc. All other mentioned trade names and/or trademarks are the property of their respective owners.