Deployment Guide

Document version 1.0

What's inside:

Deploying the BIG-IP system with IBM Security Access Manager

Welcome to the F5 deployment guide for IBM® Security Access Manager (SAM, formerly Tivoli Access manager or TAM). This guide shows how to configure the BIG-IP Local Traffic Manager (LTM) and BIG-IP Application Acceleration Manager (AAM) with IBM Security Access Manager.

When deploying IBM Identity Management, WebSEAL is a critical component of the deployment and should be designed with a high availability architecture. WebSEAL communicates with the Secure Access Manager Policy Server and provides web proxy functionality. The BIG-IP system configuration for WebSEAL is primarily focused on SSL offload, load balancing, acceleration, and security.

Deploying the BIG-IP system in front of WebSEAL completes the highly available, secure, manageable and fast architecture required by any enterprise or business.

Why F5

Using BIG-IP with SAM brings a host of benefits that complement WebSEAL's functionality.

- BIG-IP Local Traffic Manager (LTM) provides high availability for your WebSEAL environments by using health checks to direct traffic to a WebSEAL server that is available.
- BIG-IP LTM SSL offload brings step-down authentication capability to your WebSEAL deployments. By using 2048 or larger keys using ECC technology on the BIG-IP system, users can realize the strongest possible encryption while BIG-IP uses more efficient 1024 keys for communication with WebSEAL.
- BIG-IP Application Acceleration Manager (AAM) can provide content caching and intelligent browser referencing (IBR) to accelerate the user experience for the content that your WebSEAL proxies are serving. BIG-IP AAM dynamically manages expires headers, provided content caching and intelligently manages content with browsers, reducing the total number of HTTP connections between browser and server, among other acceleration features.

For more information on Security Access Manager (formerly IBM Tivoli Access Manager,) see *http://www-03.ibm.com/software/products/us/en/identity-access-manager*

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 and AAM	11.4
IBM Security Access Manager for Web	7.0

Important: *Make sure you are using the most recent version of this deployment guide, available at http://www.f5.com/pdf/deployment-guides/ibm-security-access-manager-dg.pdf.*

Prerequisites and configuration notes

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

- In order to use the BIG-IP Application Acceleration Manager (AAM), it must be fully licensed and provisioned on the BIG-IP system.
- If you are using the BIG-IP system to offload SSL or for SSL re-encryption (SSL Bridging), you must have already obtained a valid SSL certificate and key, and it is imported it onto the BIG-IP LTM system. For specific instructions on importing SSL certificates and keys, see the online help or BIG-IP system documentation, available at http://support.f5.com/kb/en-us.html
- This document is intended for the load balancing and acceleration of WebSEAL components. WebSEAL should be configured and functional on your network
- This document focuses on the availability of the proxy features of WebSEAL. It is not concerned with the load balancing or acceleration of the administration functions of WebSEAL.

Configuration example

The following simple configuration example shows the BIG-IP system with LTM and AAM modules in front of a pool of WebSEAL devices.

One of the core components of the BIG-IP system is providing high availability. In this implementation, after checking server health the BIG-IP LTM distributes user traffic to the WebSEAL server with the fewest connections. Because the user could be sent to any of the WebSEAL devices that are a part of this configuration, it is best practice that all the servers are identical. See the following section for more information on replicating the WebSEAL servers.

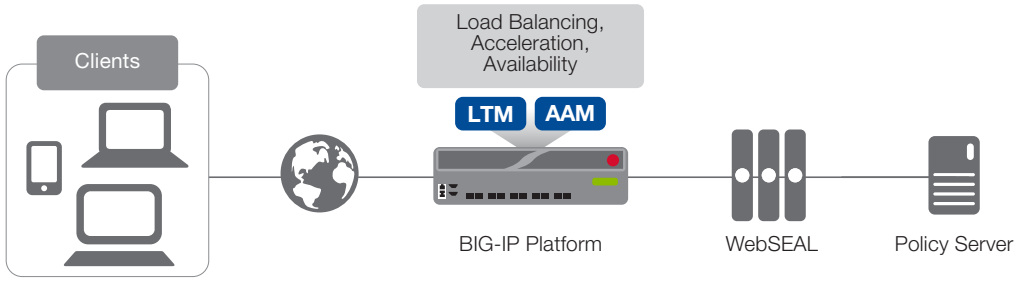

Figure 1: *Logical configuration example*

Configuring the BIG-IP System for WebSEAL

In this section, we describe how to replicate front-end WebSEAL servers, as well as how to configure the BIG-IP system for WebSEAL.

Replicating front-end WebSEAL servers

Because it is best practice that all the WebSEAL servers are identical (as described in the Configuration example section), In this procedure, we show you how to replicate front-end WebSEAL servers. For specific instructions, see the IBM documentation.

In this example, the host name of the primary WebSEAL server machine is **WS1**. The host name for the replica WebSEAL server machine is **WS2**.

To replicate the front-end WebSEAL servers

1. Install and configure WebSEAL on both the primary and replica server machines (**WS1** and **WS2** in our example).
2. Create a new object to be the root of the authorization space for both WebSEAL servers using the **pdadmin** command as shown in the For example:

   ```
   pdadmin> object create /WebSEAL/newroot "Description" 5 ispolicyattachable yes
   ```

3. Stop WebSEAL on the primary server (WS1 in our example).
4. On the primary server, change the value of the server-name stanza entry in the WebSEAL configuration file from the original host name (WS1 in our example) to **newroot**:

   ```
   [server]
   server-name = newroot
   ```

5. Restart WebSEAL on the primary server.
6. Repeat Steps 3-5 for the replica server (WS2 in our example).

The primary and replica servers now use the object /WebSEAL/newroot as the base for authorization evaluations. Either server can respond to object list and object show commands for objects located below /WebSEAL/newroot.

Configuring the BIG-IP system

Use the following table for guidance on configuring the BIG-IP LTM for WebSEAL. This table contains any non-default setting you should configure as a part of this deployment. Settings not contained in the table can be configured as applicable. For specific instructions on configuring individual objects, see the online help or product manuals.

BIG-IP Object	Non-default settings/Notes	
Health Monitor[1] (*Local Traffic > Monitors*)	***Name***	Type a unique name
	Type	**HTTP**
	Interval	**30** (recommended)
	Timeout	**91** (recommended)
Pool (*Local Traffic > Pools*)	***Name***	Type a unique name
	Health Monitor	Select the monitor you created above
	Slow Ramp Time[2]	**300**
	Load Balancing Method	Choose **Least Connections (Node)**
	Address	Type the IP Address of a WebSEAL server
	Service Port	**80** if offloading SSL, **443** if not Repeat Address and Service Port for all nodes
Optional: **BIG-IP AAM** (*Acceleration > Web Application > Application*)	***Application Name***	Type a unique name
	Policy	Select **Generic Policy - Complete**
	Requested Host	Type the domain name (host name) that might appear in HTTP requests for WebSEAL. Click **Add Host** to include additional host names.
Profiles (*Local Traffic-->Profiles*)	***HTTP*** (*Profiles-->Services*)	Name: Type a unique name Parent Profile: **http** Insert X-Forwarded-For: If you are using SNAT (recommended): **Enabled**
	HTTP Compression (*Profiles-->Services*)	Name: Type a unique name Parent Profile: **httpcompression**
	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**
	Web Acceleration[3] (*Profiles-->Protocol*)	Name: Type a unique name Parent Profile: **webacceleration** WA Applications: Enable your BIG-IP AAM application
	Client SSL (*Profiles-->SSL*)	Name: Type a unique name Parent Profile: **clientssl** Certificate and key: Select the Certificate and key you imported for this implementation
	Server SSL ***(for SSL Bridging only)*** (*Profiles-->SSL*)	Name: Type a unique name Parent Profile: If your server is using a certificate signed by a CA, select **serverssl**. If your server is using a self-signed certificate, or an older SSL cipher, select **serverssl-insecure-compatible**. Certificate and Key: Leave Certificate and Key set to **None**.

[1] To make this monitor more sophisticated, see *Adding enhanced monitoring to the implementation on page 6*
[2] You must select **Advanced** from the **Configuration** list for these options to appear
[3] Optional: Only necessary if you are deploying the BIG-IP AAM

BIG-IP Object	Non-default settings/Notes	
Virtual Server (*Main tab-->Local Traffic -->Virtual Servers*)	***Name***	Type a unique name.
	Address	Type the IP Address for this virtual server
	Service Port	**443** if offloading SSL or SSL Bridging, **80** if not.
	Protocol Profile (Client)[1]	Select the WAN optimized TCP profile you created above
	Protocol Profile (Server)[1]	Select the LAN optimized TCP profile you created above
	Web Acceleration Profile[2]	Select the Web Acceleration profile you created above
	Source Address Translation	**Auto Map** [3]
	Default Pool	Select the appropriate pool you created above

[1] You must select **Advanced** from the **Configuration** list for these options to appear
[2] Optional: Only necessary if you are deploying the BIG-IP AAM
[3] If you have a large deployment in which you expect more than 64,000 simultaneous connections per server, you must configure a SNAT Pool, with an IP address for each 64,000 simultaneous connections you expect. See the BIG-IP documentation on configuring SNAT Pools.

This completes the BIG-IP configuration.

Next steps

By completing the configuration in this guide, you have set up multiple WebSEAL servers and made sure they are identical, you have completed the BIG-IP system configuration for load balancing these WebSEAL servers, and you may have added optional SSL offload and acceleration. To ensure that you experience the maximum benefit from your new environment, we recommend the following post-configuration tasks.

Adjust DNS entries in your environment to point to the virtual IP address

In this guide you created a front-end IP address on the BIG-IP system; the virtual IP address. You should now adjust all services and users that would have been connecting directly to a WebSEAL server to this virtual address. In typical environment, this means adjusting your DNS entry to point to this virtual server IP address on the BIG-IP system.

In some environments, you may be using BIG-IP Global Traffic Manager (GTM) to distribute WebSEAL servers globally at multiple data centers. In this case, you would associate the virtual IP addresses of each WebSEAL environment with BIG-IP GTM. Please see BIG-IP documentation for further information GTM.

Adjust compression and caching settings on WebSEAL

If you are using BIG-IP AAM (Application Acceleration Manager) to cache and accelerate WebSEAL server content, you modify the WebSEAL server to further optimize the CPU, memory and disk utilization of the WebSEAL servers. We recommend disabling compression and turning off caching on WebSEAL if you use BIG-IP AAM. Please refer to WebSEAL documentation on making these adjustments.

Adding enhanced monitoring to the implementation

In this document, we describe a simple HTTP monitor that tests whether WebSEAL proxy is available. This monitor serves three purposes:

- It establishes that an individual WebSEAL server is powered on.
- It establishes that the operating system on the WebServer server is able to answer TCP requests, which also means the server has sufficient CPU cycles to allocate user time to the WebSEAL process.
- By connecting to the HTTP process, the monitor determines that the actual WebSEAL proxy is operational and has the CPU cycles necessary to answer a request.

We recommended enhancing the monitor by using the Send string and Receive string options to further exercise and test the functionality of the underlying disk subsystem and associated transport systems.

In order to modify the monitor, simply open the monitor you created in this guide, and add Send and Receive values. The Send value needs to be properly formatted HTTP and can be anything from:

```
GET / HTTP/1.0\r\n
```

to something more complicated such as

```
GET /mytesturl.html HTTP/1.1\r\nHost: myhostname.local\r\n\r\n
```

.

A Receive string is the string you would expect to receive if you executed this query from your browser. It can be something as simple as looking for a word that appears in the response, such as **WebSEAL**, or it can be a regular expression.

See the BIG-IP documentation for complete information on configuring advanced monitors.

Troubleshooting

Issue: *When sending a request through the BIG-IP Virtual Server IP address, the response does not come back, but if sending directly to a WebSEAL server, the response is received.*

Troubleshooting: In a scenario where responses do not come back when accessing a server through the BIG-IP system, the primary cause is often asymmetric routing. This means that the network connection is taking a different route back to the originating client than the one used to get to the server. This often happens when the servers do not have the BIG-IP system as their default route, which is often the case. If your WebSEAL servers do not have the BIG-IP system as their default route, make sure that you have added a SNAT (either Auto Map or a SNAT Pool) to the virtual server as described in this document.

Issue: IP Addresses in the WebSEAL logs show the BIG-IP Self IP address instead of the client's actual IP address.

Troubleshooting: When servers do not have their default route back to the BIG-IP system, SNAT must be used to avoid asymmetric routing problems which will prevent the delivery of traffic to clients. A side-effect of SNAT is that the originating IP address can be lost. In order to solve this issue, the BIG-IP system inserts an X-Forwarded-For header in the HTTP header and passes this to the server. Follow the configuration steps in this document to use a custom HTTP profile with X-Forwarded-For set to **Enabled**. Your WebSEAL proxy can be configured to log this X-Forwarded-For header or pass it on to the application server behind it.

A second side-effect of using SNAT is that IP Address-based authentication on WebSEAL will not function properly in SNAT environments. IP Address authentication uses the source IP address, which will always be the Self IP address of the BIG-IP system.

If WebSEAL IP Address based authentication is absolutely required, we recommend the WebSEAL servers be reconfigured to have a default route back to the BIG-IP system. A second option is to move the IP Address authentication to the BIG-IP system itself. This can be achieved through an iRule or through the use of the Application Policy Manager (APM) module. For more information on these options please see BIG-IP documentation or refer to F5's DevCentral site (*http://devcentral.f5.com*)

Document Revision History

Version	Description	Date
1.0	New guide	06-12-2013

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