What's inside:

Deploying the BIG-IP LTM with IBM WebSphere MQ

Welcome to the F5 Deployment Guide for IBM® WebSphere® MQ. This document provides guidance for deploying the BIG-IP Local Traffic Manager (LTM) with IBM WebSphere MQ. The BIG-IP LTM brings high availability, SSL offload, and TCP optimizations to WebSphere MQ solutions.

WebSphere MQ improves the flow of information across an organization and positions it to adjust to dynamic business requirements, reduce maintenance, integration costs, and seamlessly bridge to new technologies.

Why F5

The BIG-IP LTM brings high availability, SSL offload and TCP optimization to WebSphere MQ solutions. The primary use case addressed in this guide is placing BIG-IP LTM in front of incoming MQ queue managers for connection balancing of receiver queues. The BIG-IP LTM can also provide monitoring and high availability for transmission queues if affinity is not required.

While WebSphere MQ already provides connection balancing, utilizing BIG-IP brings a number of additional benefits.

- WebSphere MQ connection balancing is based on a static list of addresses. If one or more of these addresses are down, the WebSphere MQ client spends time trying to connect to them anyway. By using a virtual server address on the BIG-IP system as described in this deployment guide, the BIG-IP device routes each connection request directly to an available MQ instance.

- WebSphere MQ connection balancing is configured at build time using a client-channel definition table file or JMS managed object definition. By using the BIG-IP system, changes to the MQ Server list are dynamic and do not require the client application to restart or redeploy to pick up the changes.

- WebSphere MQ connection balancing is based on weighting and each connection is evaluated independently. The BIG-IP system, as deployed in this deployment guide, is using the Least Connections algorithm, which means that new connections are balanced based on the number of live connections on each node.

For information on IBM WebSphere MQ see: *http://www-01.ibm.com/software/integration/wmq/*

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	11.1 HF-2
IBM WebSphere MQ	7.1

Important: *Make sure you are using the most recent version of this deployment guide, found at http://www.f5.com/pdf/deployment-guides/ibm-websphere-mq-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:

- If you are using the BIG-IP system to offload SSL, we assume you have already obtained an SSL certificate and key, and it is installed on the BIG-IP LTM system.
- As stated in the introduction, the primary use case in this deployment guide is the BIG-IP system deployed in front of queue managers, providing load balancing and offload.
- WebSphere MQ heartbeats should be configured to a value smaller than the BIG-IP LTM TCP Idle Timeout value. We recommend 180 seconds for the BIG-IP LTM TCP Idle Timeout value (as shown in this guide) and 60 seconds for the WebSphere MQ heartbeat value. For information on configuring the WebSphere MQ heartbeats, see the IBM documentation.

Configuration example and traffic flows

Using the configuration in this guide, the BIG-IP system provides high availability directly to WebSphere Message Broker Servers. If DataPower XI50 devices are used for XML transformation in your implementation, the BIG-IP provides high availability to the DataPower devices.

The traffic flows for each of the modes, and configuration instructions are below. The setup of BIG-IP is currently identical between the two modes, but the setup of WebSphere MQ is different between the two modes.

Mode 1 - BIG-IP LTM directing traffic to WebSphere MQ

In the following diagram, the BIG-IP LTM provides intelligent traffic direction and high availability for WebSphere Message Broker servers.

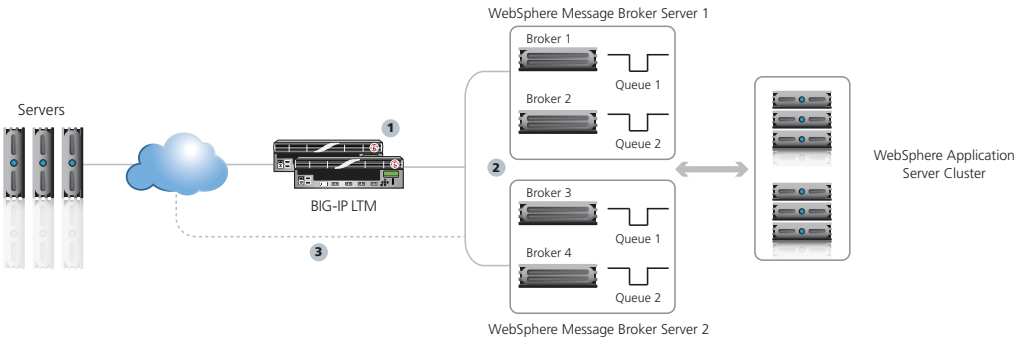

1. The BIG-IP system continually monitors the WebSphere MQ servers for health and availability.
2. The BIG-IP system accepts incoming queue messages and delivers them to the appropriate Broker server.
3. Outgoing queues may return without traversing the BIG-IP LTM.

Configuring WebSphere MQ devices for use with the BIG-IP system

To provide high availability for WebSphere MQ, you must have two or more identical WebSphere Message Broker Servers. For example, you should setup the exact same transmission queues, Queue Managers and Channels on all MQ servers, using the same TCP ports and names for all servers. For specific instructions, see the IBM documentation.

Mode 2 – Load balancing DataPower Devices

In the following diagram, the BIG-IP LTM provides intelligent traffic direction and high availability to the DataPower devices.

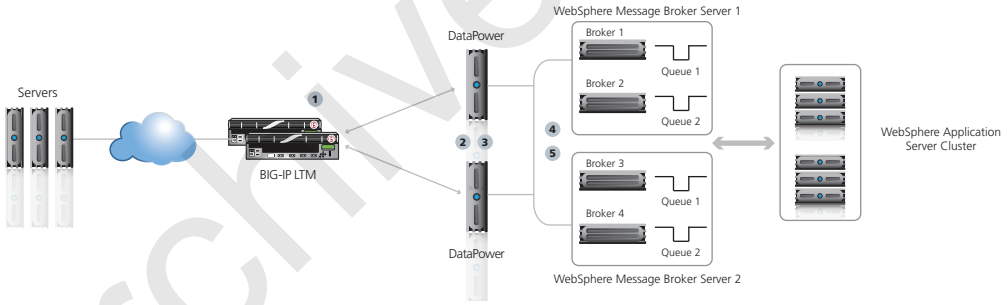

This diagram illustrates the following process:

1. The BIG-IP LTM receives all incoming requests and distributes these requests across the DataPower XI50 appliances.
2. The DataPower devices perform basic validation and threat protection on the SOAP requests. It also load balances the requests to the WebSphere Message Broker servers in the network.
3. Each broker contains two execution groups running an instance of the message flow. This results in eight instances of the same message flow. DataPower load balances across these eight endpoints.
4. The message flow writes the message to a WebSphereMQ queue.
5. The message is consumed by a MDB connected to WebSphereMQ using client bindings.

The high availability features of the topology are as follows:

- If a DataPower device becomes unavailable, traffic can be routed to an alternate device.
- If one WebSphere Message Broker server becomes unavailable, all traffic is routed to the alternate server.
- If one or more brokers becomes unavailable, all traffic is routed to the remaining brokers
- If one or more execution groups becomes unavailable, all traffic is routed to the remaining execution groups.

Relationship between MQ queue managers and BIG-IP virtual server addresses

In the following chart, we demonstrate the relationship between MQ queue managers, Port and IP information for that queue manager, and the BIG-IP virtual server. In this example, there are three queue managers, SalesQueue, OrderQueue and InventoryQueue, installed on two MQ Servers, 192.168.10.50 and 192.168.10.60. The queue managers are each mapped on a specific port on the server, in this case, 1414, 1415 and 1416. On the BIG-IP LTM, virtual servers are configured for each queue manager on the same TCP port, but in our case with external routed IP addresses. The BIG-IP LTM pool contains the two MQ servers and monitors these servers for health and availability before delivering message traffic. By separating queue managers on their own ports, persistence and grouping of messages can be managed on a more granular level, with more visibility into the health of each server.

MQ Queue Manager	Queue Manager (IP:Port)	BIG-IP virtual server
SalesQueue manager	192.168.10.50:1414 and 192.168.10.60:1414	64.0.0.1:1414
OrderQueue manager	192.168.10.50:1415 and 192.168.10.60:1415	64.0.0.1:1415
InventoryQueue manager	192.168.10.50:1416 and 192.168.10.60:1416	64.0.0.1:1416

Configuring the BIG-IP LTM

Use the following table for guidance on configuring the BIG-IP LTM for either deployment mode. This table shows the required BIG-IP configuration objects 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 described in the following table, you need to create a BIG-IP pool and virtual server for each transmission queue that is a part of this deployment. For instructions on

Important

The heartbeat value in your WebSphere MQ configuration must be less than the BIG-IP LTM Idle Timeout value in the TCP configuration. We recommend a WebSphere MQ heartbeat value of 60 seconds. See the WebSphere documentation for specific instructions on configuring the heartbeat.

*It is critical that a **tcp_half_open** monitor be used, in order to minimize impact on the WebSphere MQ server. If a full TCP monitor is used, WebSphere MQ generates a dump file and may degrade the performance of the queue manager over time.*

BIG-IP Object	Non-default settings/Notes	
Health Monitor (*Main tab-->Local Traffic -->Monitors*)	***Name***	Type a unique name
	Type	**TCP Half Open**
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 **Least Connections (Member)**
	Address	Type the IP Address of a WebSphere MQ node
	Service Port	Type the appropriate port for the channel, such as **1414**. Repeat Address and Service Port for all nodes)
	Create additional pools for each Receiver Queue be load balanced. Use the appropriate Service port for the specific Receiver Queue.	
Profiles (*Main tab-->Local Traffic -->Profiles*)	***TCP WAN*** (*Profiles-->Protocol*)	Name: Type a unique name; Parent Profile: **tcp-wan-optimized**; Idle Timeout[2]: **180**[2]
	TCP LAN (*Profiles-->Protocol*)	Name: Type a unique name; Parent Profile: **tcp-lan-optimized**; Idle Timeout[2]: **180**[2]
	Client SSL[2] (*Profiles-->SSL*)	Name: Type a unique name; Parent Profile: **clientssl**; Certificate: Select the Certificate you imported; Key: Select the associated Key
	Server SSL[3] ***(for SSL Bridging only)*** (*Profiles-->SSL*)	Name: Type a unique name; Parent Profile: If your server is using a certificate signed by a Certificate Authority, select **serverssl**. If your server is using a self-signed certificate, or an older SSL cipher, select **serverssl-insecure-compatible**.; Certificate and Key: Leave the Certificate and Key set to None.

[1] You must select **Advanced** from the **Configuration** list for these options to appear
[2] See important note above this table. The WebSphere MQ heartbeat value must be less than this Idle Timeout value.
[3] A Client SSL profile is only necessary if you want the BIG-IP system to decrypt SSL connections, typically for SSL Offload.
[4] The Server SSL profile is only necessary if you require encrypted traffic all the way to the servers. For SSL Offload (recommended), you do not need a Server SSL profile.

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	Type the same port you used for the pool, such as **1414**.
	Protocol Profile (Client)[1,4]	Select the WAN optimized TCP profile you created above
	Protocol Profile (Server)[1]	Select the LAN optimized TCP profile you created above
	SSL Profile (Client)[2]	If you created a Client SSL profile only: Select the Client SSL profile you created above
	SSL Profile (Server)[3]	If you created a Server SSL profile for SSL Bridging only: Select the Server SSL profile you created above.
	SNAT Pool	**Auto Map**
	Default Pool	Select the appropriate pool you created above
	Create additional virtual servers for each pool you created above. Make sure to use the appropriate Service Port, and select the appropriate Pool. You can use the same profiles.	

[1] You must select **Advanced** from the **Configuration** list for these options to appear
[2] A Client SSL profile is only necessary if you want the BIG-IP system to decrypt SSL connections, typically for SSL Offload.
[3] The Server SSL profile is only necessary if you require encrypted traffic all the way to the servers. For SSL Offload (recommended), you do not need a Server SSL profile.
[4] If the majority of your clients are connecting via a LAN, use the LAN optimized profile you created.

This completes the BIG-IP LTM configuration.

Next Steps

Now that you've completed the BIG-IP system configuration for IBM WebSphere MQ, here are some examples of what to do next.

Adjust your DNS settings to point to the BIG-IP system

After the configuration is completed, your DNS configuration should be adjusted to point to the BIG-IP virtual server for WebSphere MQ.

Advertise new Queue IP addresses to Messaging systems.

You must advertise your new Queue IP addresses to your Messaging Systems. Be sure to update your transmission queues to point to the BIG-IP LTM virtual IP address or the DNS name you have created for this address.

If you do not advertise the IP addresses, traffic is sent directly to the broker servers and not the high availability system you have just created.

Make sure the BIG-IP TCP Idle Timeout is configured properly

If you notice your WebSphere Queues are timing out, check to make sure you the WebSphere MQ heartbeat are set to a value that is smaller than the BIG-IP TCP Idle Timeout value, as described in this guide on *page 5* .

Document Revision History

Version	Description	Date
1.0	New guide	06-13-2012
1.1	- Added new content to the Why F5 section on the first page - Changed references to "MQ queues" to "MQ queue managers" - Changed the BIG-IP health monitor from TCP to TCP Half Open and added to the important note before the configuration table about why the TCP Half Open monitor is necessary.	03-13-2013
1.2	- Modified the parent profiles for the TCP profiles from **wom-tcp-lan-optimized** and **wom-tcp-wan-optmized** to **tcp-lan-optimized** and **tcp-wan-optimized**. - Added a note to the Protocol Profile (Client) setting on the virtual server stating if most clients are connected via a LAN, use the tcp-lan-optimized profile you created.	02-21-2014

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