



Deploying the BIG-IP System v11 with Diameter Servers

What's inside:

- 2 Prerequisites and configuration notes
- 2 Traffic flow
- 4 Preparation Worksheet
- 5 Configuring the BIG-IP iApp for Diameter
- 8 Next steps
- 9 Additional information: Why Diameter support is valuable

Welcome to the F5 deployment guide for Diameter traffic management. This guide provides step-by-step procedures for configuring the BIG-IP system version 11 for load balancing and intelligent traffic management for the Diameter protocol. BIG-IP version 11.0 introduces iApp™ Application templates, an extremely easy way to accurately configure the BIG-IP system for your Diameter servers.

The Diameter base protocol is intended to provide an Authentication, Authorization and Accounting (AAA) framework for applications such as network access or IP mobility. Diameter is also intended to work in both local Authentication, Authorization & Accounting and roaming situations (this is an excerpt from RFC3588; for more information, see the complete RFC: <http://www.ietf.org/rfc/rfc3588.txt>).

Why F5?

The BIG-IP system provides a number of ways to optimize and scale Diameter deployments. By supporting message-based load balancing, the BIG-IP LTM can act as a proxy that will de-multiplex each request from the client to multiple servers and improve overall performance and scalability. For more information, see *Additional information: Why Diameter support is valuable on page 9*

To provide feedback on this deployment guide or other F5 solution documents, contact us at solutionsfeedback@f5.com.

Products and versions tested

Product	Version
BIG-IP LTM	v11

Important: Make sure you are using the most recent version of this deployment guide, available at <http://www.f5.com/pdf/deployment-guides/diameter-iapp-dg.pdf>.

What is F5 iApp™?

New to BIG-IP version 11, F5 iApp is a powerful new set of features in the BIG-IP system that provides a new way to architect application delivery in the data center, and it includes a holistic, application-centric view of how applications are managed and delivered inside, outside, and beyond the data center. The iApp template for Diameter acts as the single-point interface for building, managing, and monitoring your Diameter deployment.

For more information on iApp, see the White Paper *F5 iApp: Moving Application Delivery Beyond the Network*: <http://www.f5.com/pdf/white-papers/f5-iapp-wp.pdf>.

Document Version

1.0

Prerequisites and configuration notes

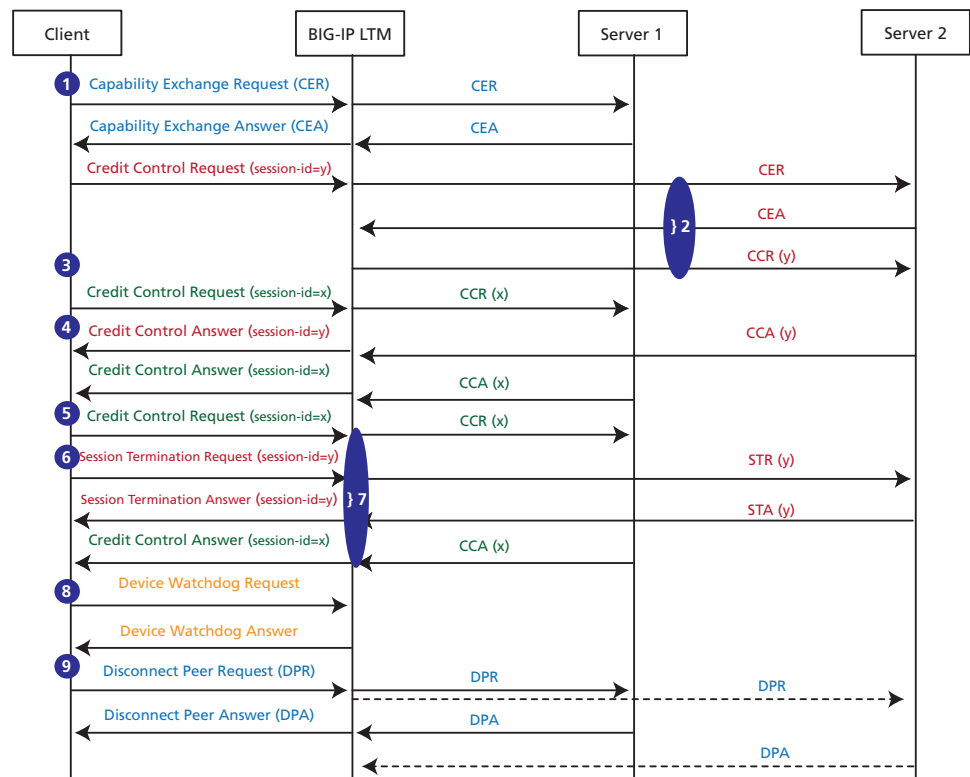
The following are general prerequisites and configuration notes for this guide:

- For this deployment guide, the BIG-IP LTM system **must** be running version 11.0 or later. If you are using a previous version of the BIG-IP LTM system, see the Deployment Guide index on F5.com. The configuration described in this guide does not apply to previous versions.
- This deployment guide provides guidance for using the iApp for Diameter found in version 11.0 and later. For advanced users extremely familiar with the BIG-IP, there is a manual configuration table at the end of this guide. However, we strongly recommend using the iApp template.
- If you are using the BIG-IP system to offload TLS processing, we assume you have already obtained a valid certificate and key, and it is installed on the BIG-IP LTM system.

Traffic flow

The following diagram contains an example of the traffic flow for a load balanced Diameter implementation.

This traffic flow diagram is written with the following assumptions: this is the first time the LTM has delivered traffic to the server and there are no persistent records; the round robin load balancing algorithm is used, and the persistent key is session-id AVP.



1. The Diameter Handshake or Capability Exchange Request is initiated and load balanced by the BIG-IP LTM to server 1.
2. The next Diameter request has session-id = y. It has never been seen before, so the BIG-IP LTM sends the request to server 2. However, prior to sending the request, the BIG-IP LTM

performs the Diameter Handshake with the server. The BIG-IP LTM uses the session-id AVP as a persistent key. This request has session-id = y.

3. This request has session-id = x. It has never been seen before, so the BIG-IP LTM sends the request to server 1.
4. This request has session-id = y. The BIG-IP LTM persists this request to server 2.
5. This request has session-id = x. The BIG-IP LTM persists this request to server 1.
6. This request has session-id = y. The BIG-IP LTM persists this request to server 2.
7. The Diameter protocol is asynchronous. The response from the server can come any time and in any order. Note STA(y) is sent back to the client prior to CCA(x).
8. When the client or server sends a Device Watchdog Request, the BIG-IP LTM responds with a Device Watchdog Answer (the BIG-IP LTM does not forward DWR).
9. When the client sends the Disconnect Peer Request, the LTM broadcasts the request to all servers.
If the server sends the Disconnect Peer Request, the LTM responds with a Disconnect Peer Answer to that particular server only. The connections to the client and other servers are not affected.

Preparation Worksheet

In order to use the iApp for Diameter, you need to gather some information, such as server IP addresses and domain information. Use the following worksheet to gather the information you will need while running the template. The worksheet does not contain every question in the template, but rather include the information that is helpful to have in advance. More information on specific template questions can be found on the individual pages.

You might find it useful to print this table and then enter the information.

➤ **Note:** *Although we show space for 10 pool members, you may have more or fewer members in each pool.*

IP Addresses/FQDN	Port	TLS Offload	Protocol	Pool Members	Sync/Failover Groups*
<p>IP address you will use for the LTM virtual server:</p> <p>FQDN that will resolve to the virtual server address:</p>	<p>What port will you use for the diameter virtual server?</p> <p>The default port is 3868</p>	<p><i>Offloading TLS?</i> Yes No</p> <p>If offloading TLS, import a certificate and key into the BIG-IP LTM before running the template.</p> <p>Certificate:</p> <p>Key:</p>	<p>Are you using TCP or SCTP</p>	<p>Diameter server IP addresses:</p> <p>1:</p> <p>2:</p> <p>3:</p> <p>4:</p> <p>5:</p> <p>6:</p> <p>7:</p> <p>8:</p> <p>9:</p> <p>10:</p> <p>Port used by Diameter (3868 is the default):</p>	<p>If using the Advanced feature of Sync/Failover Groups, you must already have a Device Group and a Traffic Group</p> <p>Device Group name:</p> <p>Traffic Group name:</p>
<p>TCP request queuing*</p> <p>If using TCP request queuing, you should know the queue length and timeout, as well as the connection limit for the node.</p> <p>Request queue length:</p> <p>Timeout:</p> <p>Node Connection limit:</p>					

* *Optional*

Configuring the BIG-IP iApp for Diameter

Use the following guidance to help you configure the BIG-IP system for Diameter using the BIG-IP iApp template.

Getting Started with the iApp for Diameter

To begin the Diameter iApp Template, use the following procedure.

1. Log on to the BIG-IP system.
2. On the Main tab, expand **iApp**, and then click **Application Services**.
3. Click **Create**. The Template Selection page opens.
4. In the **Name** box, type a name. In our example, we use **Diameter_**.
5. From the **Template** list, select **f5.diameter**.
The Diameter template opens.

Advanced options

If you select **Advanced** from the **Template Selection** list, you see Sync and Failover options for the application. This feature, new to v11, is a part of the Device Management configuration. This functionality extends the existing High Availability infrastructure and allows for clustering, granular control of configuration synchronization and granular control of failover. For more information on Device Management, see the Online Help or product documentation.

1. **Configure Sync/Failover?**
If you want to configure the Application for Sync or failover groups, select **Yes** from the list.
 - a. **Device Group**
If you select Yes from the question above, the Device Group and Traffic Group options appear. If necessary, uncheck the Device Group box and then select the appropriate Device Group from the list.
 - b. **Traffic Group**
If necessary, uncheck the Traffic Group box and then select the appropriate Traffic Group from the list.

Virtual Server Questions

The next section of the template asks questions about the BIG-IP virtual server. A virtual server is a traffic-management object on the BIG-IP system that is represented by an IP address and a service. Clients send traffic to a virtual server, which then directs the traffic according to your configuration instructions.

1. **IP address for the virtual server**
This is the address clients use to access the Diameter servers (or a FQDN will resolve to this address). You need an available IP address to use here.
2. **Port**
If you are using a port other than the default port for Diameter (3868), type it here.
3. **Routes or secure network address translation**
If the Diameter devices do not have a route back for clients through the BIG-IP, i.e. if they do not use the BIG-IP as the default gateway, the BIG-IP uses Secure Network Address Translation (SNAT) Automap (exception in #3) to translate the client's source address to an address configured on the BIG-IP. The servers then use this new source address as the destination address when responding to traffic originating through the BIG-IP.

If you indicate the Diameter devices do have a route back to the clients through the BIG-IP, the BIG-IP does not translate the client's source address; in this case, you must make sure that the BIG-IP is configured as the gateway to the client networks (usually the default gateway) on the Diameter devices.

We recommend choosing **No** from the list because it is secure and does not require you to configure routing manually.

If you are configuring your BIG-IP LTM in a "one-armed" configuration with your Diameter devices -- where the BIG-IP virtual server(s) and the Diameter devices have IP addresses on the same subnet -- you must choose **No**.

4. **More than 64,000 simultaneous connections**

If you do not expect more than 64,000 simultaneous connections, leave this answer set to **No** and continue with the next section.

If you have a large deployment and expect more than 64,000 connections at one time, the iApp creates a SNAT Pool instead of using SNAT Automap. With a SNAT Pool, you need one IP address for each 64,000 connections you expect. Select **Yes** from the list. A new row appears with an IP address field. In the **Address** box, type an IP address and then click **Add**. Repeat for any additional IP addresses.

5. **Protocol**

From the list, select whether your implementation uses the TCP or SCTP protocol.

6. **TLS Offload**

Before running the template you should have already imported a certificate and key onto the BIG-IP system.

For information on SSL certificates on the BIG-IP system, see the online help or the *Managing SSL Certificates for Local Traffic* chapter in the **Configuration Guide for BIG-IP Local Traffic Manager** available at <http://support.f5.com/kb/en-us.html>.

To configure the BIG-IP to offload TLS, select **Yes** from the list.

7. **Certificate**

Select the certificate for you imported for the Diameter devices from the certificate list.

8. **Key**

Select the associated key from the list.

Server Pool, Load Balancing, and Service Monitor questions

In this section, you add the Diameter servers, and configure the health monitor and pool.

1. **New Pool**

Choose **Create New Pool** unless you have already made a pool on the LTM for the Diameter devices.

2. **Load balancing method**

Choose any of the load balancing methods from the list. We use the default, **Round Robin**.

3. **Address/Port**

Type the IP Address and Port for each Diameter device (the default port is 3868). You can optionally add a Connection Limit (see note on the left). Click **Add** to add additional servers to the pool.

➔ **Important**

If you enable TCP Request Queuing, you must have a Connection Limit set on at least one of the nodes when configuring the Address/Port

4. **TCP Request Queuing**

TCP request queuing provides the ability to queue connection requests that exceed the capacity of connections for a pool as determined by the connection limit. Consequently, instead of dropping connection requests that exceed the capacity of a pool, TCP request queuing enables those connection requests to reside within a queue in accordance with defined conditions until capacity becomes available. For more information on TCP Request Queuing, see the *New Features Guide for BIG-IP Version 11*, available on Ask F5.

If you want the BIG-IP to queue TCP requests, select **Yes** from the list. Additional options appear.

- a. Type a queue length in the box. Leave the default of 0 for unlimited.
- b. Type a number of milliseconds for the timeout value.

5. **Health Monitor Interval**

The BIG-IP system creates a Diameter-specific health monitor for this configuration. Specify a time out value. We recommend the default, 30 seconds.

Protocol Optimization Questions

In this section, you configure protocol optimizations.

1. **Server-initiated messages**

In most Diameter traffic management deployments, clients send requests and servers reply with responses. However, the Diameter protocol is also designed to support server-initiated messages. There are some diameter applications/deployments that not only clients that send requests, servers may send the requests as well.

If your Diameter implementation supports server-initiated messages, select **Yes** from the list.

2. **Persistence**

If persistence is required in your implementation, select **Yes** from the list. Another row appears with the two options for persistence, **Universal** and **CARP Hash**.

3. **Persistence Type**

If you selected Yes from the Persistence list, you must choose between the Universal and CARP Hash persistence methods.

- **Universal:** If you select Universal, the BIG-IP system creates a Universal persistence profile that is used by the Diameter profile created by the template.
- **CARP Hash:** You can choose to configure CARP hash persistence for increased performance and scalability. This method is suitable for organizations who want the diameter messages to be persisted for very long time periods.

The CARP hash persistence method does not store the persistence records, so it reduces delay and the amount of processing power used by persistence look ups. This is also reduces processing power and network traffic in case the high availability deployment and persistence records have to be mirrored. Because CARP hash does not store the persistent records, so there is no need to mirror any persistent information across HA units.

The disadvantage of using CARP hash persistence compared to using universal persistence is when a new server is added (or comes back online after the health monitor detects it was down). The hash result changes and some requests from clients which belong to active sessions may be forwarded to the wrong server. However, in theory, if a single server is added, only

around $1/N$ (where N is total number of servers in the pool) of requests may have their hash result change.

If CARP hash persistence is used, we recommend that the new server should be added to the pool during a maintenance period or a time with the least amount of traffic.

For more information on the CARP hash algorithm, see <https://support.f5.com/kb/en-us/solutions/public/11000/300/sol11362.html> or the BIG-IP documentation.

Finished

Review the answers to your questions. When you are satisfied, click the **Finished** button. The BIG-IP system creates the relevant objects for the Diameter implementation.

Next steps

After completing the iApp Template, the BIG-IP Application Services page opens for the Diameter service you just created. To see the list of all the configuration objects created to support Diameter, on the Menu bar, click **Components**. The complete list of all Diameter related objects opens. You can click individual objects to see the settings.

Once the objects have been created, you are ready to use the new deployment.

Modifying DNS settings to use the BIG-IP virtual server address

Before sending traffic to the BIG-IP system, your DNS administrator may need to modify any DNS entries for the Diameter implementation to point to the BIG-IP system's virtual server address.

Modifying the iApp configuration

The iApp application service you just created can be quickly and easily modified if you find it necessary to make changes to the configuration. The Strict Updates feature of the iApp prevents users from manually modifying the iApp configuration (Strict Updates can be turned off, but use extreme caution). As a safer option, the iApp allows you to re-enter the template, make changes, and then update the template. The modifications are automatically made to any of the associated objects.

To modify the configuration

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. Click the name of your Diameter Application Service from the list.
3. On the Menu bar, click **Reconfigure**.
4. Make the necessary modifications to the template.
5. Click the **Finished** button.

Viewing statistics

You can easily view a number of different statistics on the BIG-IP system related to the Diameter configuration objects created by the iApp template.

Object-level statistics

Use the following procedure to view object-level statistics.

To view object-level statistics

1. On the Main tab, expand **Overview**, and then click **Statistics**.
2. From the **Statistics Type** menu, you can select **Virtual Servers** to see statistics related to the virtual servers.
3. You can also choose **Pools** or **Nodes** to get a closer look at the traffic.
4. To see Networking statistics in a graphical format, click **Dashboard**.

For more information on viewing statistics on the BIG-IP system, see the online help or product documentation.

Additional information: Why Diameter support is valuable

The ability to support diameter traffic management is extremely valuable. Consider the following example. In a typical load balancing situation, there are X number of clients and Y number of servers. If all clients generate one connection, there are X connections total. The BIG-IP LTM may balance X/Y connections to each server (which may be called connection-based load balancing). However, in a Diameter environment, the number of clients is likely to be small (X may even lower than Y) and that implies low number of connections (X). Moreover, each connection is long-lived, which provides few opportunities to load balance Diameter traffic on a per-connection basis.

Multiple sessions may be established within the one transport connection. Diameter keeps transport connections (TCP/SCTP) alive and reuses them for many Diameter sessions. Each Diameter session may contain multiple messages. Diameter protocol is asynchronous, or in other words, a client can send a new request without waiting for response for the previous request. The Server can send a response in any order, and can also send request.

In a high load environment, there is a need for per-message load balancing or message-based load balancing instead of connection-based load balancing. Imagine there is one transport connection between each client (NAS) and server (Diameter host server). The work required for a Diameter server to generate a response is significantly higher than the work required for a client to generate a request. Because of this, the Diameter server becomes a performance bottleneck for AAA requests from a single client. All requests from a particular client which using the same transport connection are served by only one server. By supporting message-based load balancing, the BIG-IP LTM may act as a proxy that will de-multiplex each request from the client to multiple servers and improve overall performance and scalability.

Using CARP Hash persistence

You can choose to configure CARP hash persistence for increased performance and scalability. This method is suitable for organizations who want the diameter messages to be persisted for a very long time periods.

The CARP hash persistence method does not store the persistence records, so it reduces delay and the amount of processing power used by persistence look ups. This is also reduces processing power and network traffic in case the high availability deployment and persistence records have to be mirrored. Because CARP hash does not store the persistent records, so there is no need to mirror any persistent information across HA units.

The disadvantage of using CARP hash persistence compared to using universal persistence is when a new server is added (or comes back online after the health monitor detects it was down). The hash result changes and some requests from clients which belong to active sessions may be forwarded to the wrong server. However, in theory, if a single server is added, only around $1/N$ (where N is total number of servers in the pool) of requests may have their hash result change.

If CARP hash persistence is used, we recommend that the new server should be added to the pool during a maintenance period or a time with the least amount of traffic. We also recommend you enable the manual resume option for health monitor.

For more information on the CARP hash algorithm, see <https://support.f5.com/kb/en-us/solutions/public/11000/300/sol11362.html> or the BIG-IP documentation.

Appendix: Manual configuration table

We strongly recommend using the iApp template to configure the BIG-IP system for Diameter. Advanced users extremely familiar with the BIG-IP system can use the following table to manually configure the BIG-IP system.

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

BIG-IP LTM Object	Non-default settings/Notes			
Health Monitor <i>(Main tab-->Local Traffic -->Monitors)</i>	Name	Type a unique name		
	Type	Diameter		
	Interval	30 (recommended)		
	Timeout	91 (recommended)		
Pool <i>(Main tab-->Local Traffic -->Pools)</i>	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 use the default Round Robin		
	Address	Type the IP Address of the Diameter nodes		
	Service Port	80 (click Add to repeat Address and Service Port for all nodes)		
Profiles <i>(Main tab-->Local Traffic -->Profiles)</i>	Protocol² <i>(Profiles-->Protocol)</i> Choose either TCP or SCTP	TCP	Name	Type a unique name
		Parent Profile	TCP	
	SCTP	Name	Type a unique name	
		Parent Profile	SCTP	
	Diameter <i>(Profiles-->Protocol)</i>	Name	Type a unique name	
		Parent Profile	Diameter	
		Persist Attribute ³	None (only if you are not using persistence)	
	Persistence <i>(Profiles-->Persistence)</i> If persistence is required, chose either Universal or CARP Hash	Universal	Name	Type a unique name
		Parent Profile	Persistence Type	Universal
		CARP Hash ⁴	Name	Type a unique name
			Parent Profile	Hash
	Hash Algorithm Hash Algorithm	Hash Algorithm	CARP	
mblb <i>(tmsh command line)</i>				
See <i>Creating the Message Based Load Balancing profile on page 12</i> for instructions.				
Client SSL⁵ <i>(Profiles-->SSL)</i>	Name	Type a unique name		
	Parent Profile	clientssl		
	Certificate	Select the Certificate you imported		
	Key	Select the associated Key		
iRule⁶ <i>(Main tab-->Local Traffic -->Profiles)</i>	Name	Type a unique name.		
	Definition	<pre> when LB_FAILED { if { [active_members [LB::server pool]] > 0 } { after 100 LB::reselect pool [LB::server pool] } } </pre>		

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

² Choose either TCP or SCTP, depending on the protocol you are using

³ Only modify the Persist Attribute field if you do not require persistence in your deployment

⁴ See *Using CARP Hash persistence on page 9* for more information on CARP Hash.

⁵ Only necessary if you are using the BIG-IP LTM to offload TLS

⁶ This iRule prevents client connections from being reset between monitor intervals.

Configuration table continued

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
	Service Port	3868
	Protocol Profile (client)	Select the TCP or SCTP profile you created above
	Diameter Profile	Select the Diameter profile you created above
	SSL Profile (client)¹	Select the Client SSL profile you created above
	SNAT Pool ²	Automap (optional; see footnote ²)
	iRule	Enable the iRule you created above
	Default Pool	Select the pool you created above
	Persistence Profile	Select the Persistence profile you created, if applicable

¹ Only necessary if offloading TLS

² If want to use SNAT, and you have a large deployment expecting more than 64,000 simultaneous connections, 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.

Creating the Message Based Load Balancing profile

In this section, we create the Message based load balancing (MBLB) profile. This must be done from the TMSH command line, after creating the virtual server as described above.

To add the profile using the tmsh shell

1. On the BIG-IP system, start a console session.
2. Type a user name and password, and then press Enter.
3. To add the profile to the virtual server, use the following syntax:

```
tmsh modify ltm virtual <virtual server name> profiles add { mblb }
```

In our example, we type:

```
tmsh modify ltm virtual diameter-virtual profiles add { mblb }
```

4. To save the changes, type the following command:

```
tmsh save sys config
```

Document Revision History

Version	Description
1.0	New Version

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

**F5 Networks,
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

