Deployment Guide
Document Version 1.2

Deploying the BIG-IP System with Microsoft Lync Server 2010 and 2013 for Site Resiliency

Welcome to the F5 and Microsoft® Lync Server® Deployment Guide for site resiliency. This guide contains instructions on configuring F5 Global Traffic Manager (GTM) and BIG-IP Local Traffic Manager (LTM) modules to support site resiliency for Microsoft Lync Server 2010 and 2013.

For more information on the F5 devices in this guide, see *http://www.f5.com/products/big-ip/*.

You can also visit the Microsoft page of F5's online developer community, DevCentral, for Microsoft forums, solutions, blogs and more: *http://devcentral.f5.com/Microsoft/*.

Why F5

Microsoft Lync Server supports DNS Load Balancing for distributing non-HTTP based connections to Lync Front End, Edge, and Mediation Servers. However, there are advantages to using the BIG-IP system to load balance these connections in certain scenarios:

- You are using legacy clients or servers (Office Communication Server (OCS), pre-2010 SP1 Exchange Unified Messaging).

- You are federated with an organization using legacy clients or servers (see examples in the previous bullet) or public IM services/XMPP.

- You want to use intelligent logic (such as pool member status, least connection load balancing, and more) to make the initial client connection to the Front End, Edge, or Mediation pool.

- You want to use the F5 iApp template to deploy Lync Server rapidly.

- You use F5 GTM in your environment and would like to take advantage of F5's Lync Site Resiliency solution.

Advantages

- The BIG-IP system does not interfere with Lync client registration/deregistration.

- Disabling BIG-IP system pool members is effectively the same as (and can be used in conjunction with) Lync Server Draining.

- The *BIG-IP is a certified ICSA firewall*, providing security over deploying your Lync Servers directly on the Internet.

- Even if you are using DNS Load Balancing, Microsoft requires hardware load balancing for Lync Web Services.

Microsoft®
Partner Network

Products and versions tested

Product	Versions
BIG-IP system	11.0, 11.0.1, 11.1, 11.2, 11.3, 11.4
Microsoft Lync Server	2010 and 2013

Important: *Make sure you are using the most recent version of this deployment guide, available at http://www.f5.com/pdf/deployment-guides/lync-2010-site-resiliency-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:

➤ For this deployment guide, the BIG-IP system **must** be running version 11.0 or later. The configuration described in this guide does not apply to previous versions.

➤ This document is written with the assumption that you are familiar with both F5 devices and Microsoft Lync Server. For more information on configuring these devices, consult the appropriate documentation.

➤ This guide assumes that you have configured the Lync Server Simple URLs in the following format:

» http://meet.mydomain.com

» http://dialin.mydomain.com

➤ We strongly recommend performing the BIG-IP LTM portion of this configuration using the downloadable iApp template, as described in *Downloading, importing, and running the Lync Server iApp on page 5.*

➤ You must have your internal name resolution configured properly. For example, every Lync Edge server needs to be able to resolve every Front End and Director server, and vice versa. Additionally, the FQDNs of all Lync Server pools (Edge, Front End, and Director) must resolve to the IP addresses of the corresponding BIG-IP LTM virtual servers.

➤ You must have a minimum of two BIG-IP GTM devices to perform the configuration in this guide.

Configuration example

In this guide, we configure the BIG-IP GTM with two data centers. Lync Server 2010 is configured with a "stretched" VLAN where all Lync Front End and Director Servers reside on the same layer 2 network, while Lync Server 2013 is configured with paired Front End pools for each data center (for more information on pairing pools, see *http://technet.microsoft.com/en-us/library/jj205293.aspx*).

BIG-IP GTM uses topology load balancing to direct internal clients to Lync Front End Server resources, and external clients to Edge Server and Reverse Proxy resources. The BIG-IP system monitors the availability of the internal Lync services, and marks the external services down should the health check fail. Resiliency for file shares and Microsoft SQL Server instances are beyond the scope of this document.

Logical configuration examples with traffic flow callouts for internal and external clients are on the following pages.
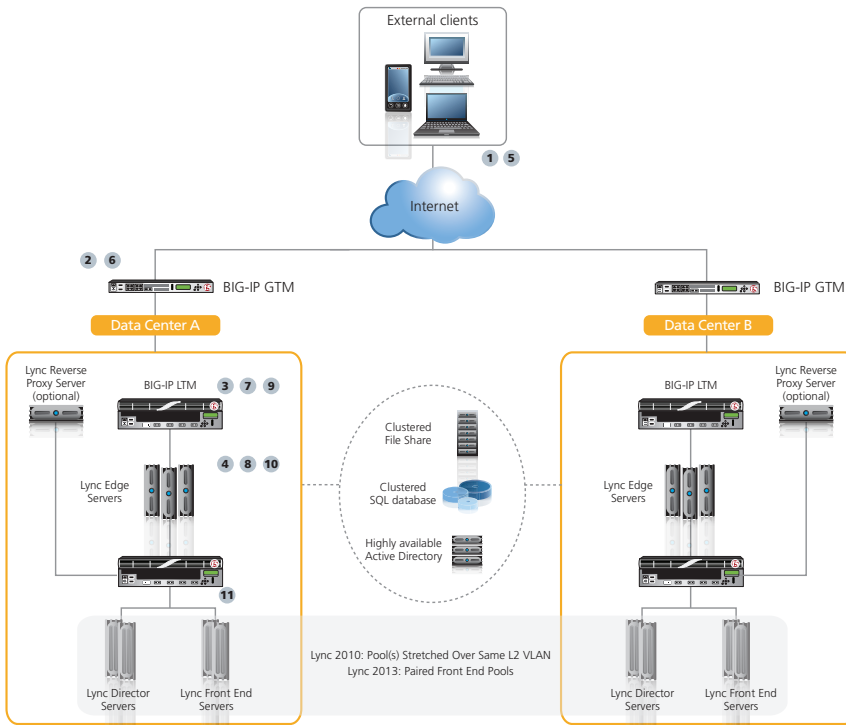
Figure 1: *Logical configuration example: External Clients*

The following describes the traffic flow for external clients:

1. An external client issues a DNS request for sip.example.com.

2. The BIG-IP GTM receives the request, and based on Topology or other intelligent load balancing calculations, responds with the external IP address of the Lync Access Service LTM virtual server in Data Center A.

3. The client connects to the Access Service virtual server on the BIG-IP LTM. The BIG-IP LTM sends the request to the Edge Server running the Access Service with the fewest connections.

4. The Access Service responds with External Web Services (Front End or Director) and Edge Services FQDNs.

5. The client issues a DNS request for Web Services and Edge Services FQDNs.

6. The BIG-IP GTM responds with the IP addresses of Reverse Proxy and Edge BIG-IP LTM virtual servers.

7. The client connects to the BIG-IP LTM Reverse Proxy for External Web Services and Edge Service virtual servers when appropriate.

8. The Edge Servers establish direct connection with the client using a local public IP address.

9. The external LTM forwards the Reverse Proxy traffic to the Front End or Director virtual servers on internal LTM.

10. Edge Servers forward Access/Conferencing/AV traffic to Lync Front End or Director Servers where applicable.

11. The internal BIG-IP LTM sends requests for simple URLs, Lync Mobility, and the Lync Address Book Service to Front End or Director Servers with the fewest connections.
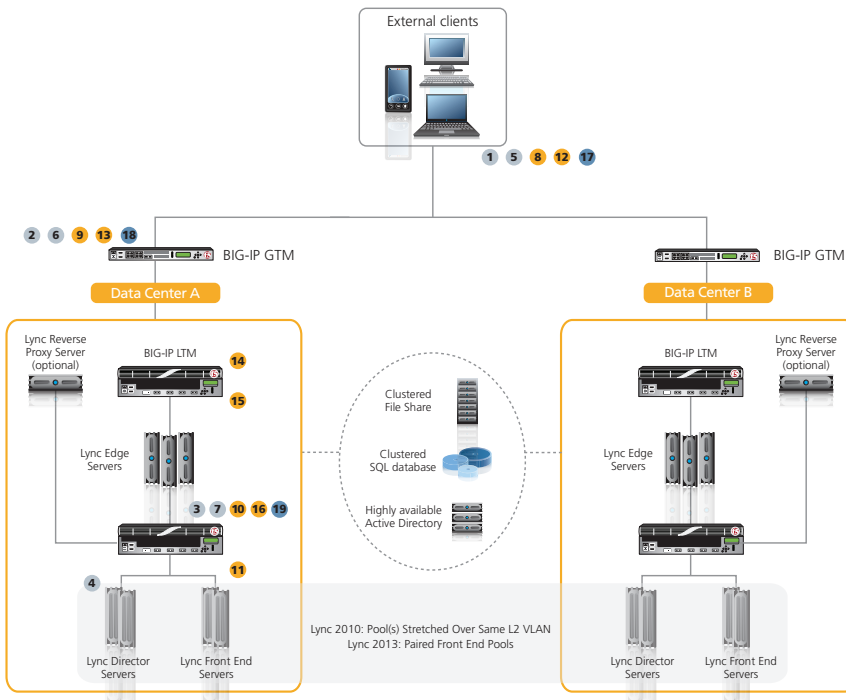
Figure 2: *Logical configuration example: Internal Clients*

The following is the traffic flow for internal clients, and is broken up into three sections. Each is represented in a different color.

Internal Clients connecting to Lync services and each other internally:

1. An internal client issues a DNS request for director.example.com.
2. The BIG-IP GTM responds with the internal IP address of the Director virtual server on the internal BIG-IP LTM.
3. The BIG-IP LTM sends the request to the Director Server with the fewest connections.
4. The Director server responds to the client with Internal Web Services FQDN (frontend.example.com).
5. The client issues a DNS request for frontend.example.com.
6. The BIG-IP GTM responds with the IP address of the Front End virtual server on the LTM.
7. The client connects to the Front End virtual server on the internal BIG-IP LTM.

Internal mobile clients connecting to the Lync Mobility service:

8. Internal DNS request for lyndiscoverinternal.example.com (internal Lync Mobility URL).
9. The BIG-IP GTM responds with the IP address of the Front End or Director virtual server on the BIG-IP LTM.
10. The client connects to the Front End or Director virtual server.
11. If external Lync Mobility is enabled, Front End or Director responds with lyncdiscover.example.com (external Lync Mobility URL).
12. The client issues a DNS request for lyncdiscover.example.com.
13. The GTM responds with IP address of the external Reverse Proxy virtual server on the External BIG-IP LTM, or directly to the server.
14. The client connects to external Reverse Proxy virtual server or server.
15. The external BIG-IP LTM, or reverse proxy server, forwards traffic to the internal LTM Reverse Proxy virtual servers.
16. The internal BIG-IP LTM forwards traffic to Front End or Director Servers.

Internal clients initiating connections to external resources (remote clients or federated IM):

17. The internal client issues DNS request for internal Edge FQDN (edge.example.com).
18. The BIG-IP GTM responds with the IP address of the Edge Internal virtual server.
19. The client connects to Edge Internal virtual server.

Configuring the BIG-IP LTM using the Lync iApp

The first task is to configure the BIG-IP LTM for Lync using the downloadable iApp available on DevCentral. For instructions on configuring the iApp, see the deployment guide at
http://www.f5.com/pdf/deployment-guides/microsoft-lync-iapp-dg.pdf.

Downloading, importing, and running the Lync Server iApp

Download the latest iApp for Microsoft Lync from DevCentral and import it onto the BIG-IP LTM. This version contains numerous enhancements and fixes to the version shipping with the product.

To download and import the iApp from DevCentral

1. Open a web browser and go to
    http://devcentral.f5.com/wiki/iApp.Microsoft-Lync-Server-2010-Updated-iApp.ashx

2. Download the **microsoft_lync_server<latest-version>.zip** file to a location accessible from your BIG-IP system.

    You must download the file, and not copy and paste the contents. F5 has discovered the copy paste operation does not work reliably.

3. Extract (unzip) the **microsoft_lync_server<latest-version>.tmpl** file.

4. Log on to the BIG-IP system web-based Configuration utility.

5. On the Main tab, expand **iApp**, and then click **Templates**.

6. Click the **Import** button on the right side of the screen.

7. Click a check in the **Overwrite Existing Templates** box.

8. Click the **Browse** button, and then browse to the location you saved the iApp file.

9. Click the **Upload** button. The iApp is now available for use.

10. On the Main tab, under **iApp**, click **Application Services**.

11. Click **Create**. The Template Selection page opens.

12. In the **Name** box, type a name. In our example, we use **Lync-server_**.

13. From the **Template** list, select **f5.microsoft_lync_server.<latest version>**.
    The new Microsoft Lync Server template opens.

14. Configure the iApp as applicable for your configuration, however, you *must* answer **Yes** to the question: "Would you like to monitor the health of the internal SIP virtual servers, and mark the Edge Access service down if monitoring fails?"

Post iApp configuration

After completing the BIG-IP LTM iApp configuration, run the Lync Topology Builder and then publish the topology. Once the topology has been published verify the availability of the Lync services through the BIG-IP LTM before continuing with the GTM configuration.

Configuring the BIG-IP GTM

Use the following tables for guidance on configuring the BIG-IP GTM. These tables contain a list of BIG-IP configuration objects you should configure as a part of this deployment. The options for the individual objects depend on your configuration. The settings we show in the following tables are provided as an example. For specific instructions on configuring individual objects, see the online help or product manuals.

F5 recommends configuring GTM Links and associating them with Data Centers. When a GTM Link associated with a Data Center is marked down, GTM no longer sends responses for resources located in that Data Center. For more information about configuring GTM Links, see the BIG-IP GTM documentation.

BIG-IP GTM Object	Non-default settings/Notes	
Data Center (*Global Traffic -->Data Centers*)	*Name*	Type a unique name. All other fields are optional.
	Important: Create a GTM Data Center for each location in your Lync environment	
DNS Profile (*Local Traffic -->Profiles -->Services-->DNS*)	*Name*	Type a unique name.
	Use BIND Server on BIG-IP	**Disabled**.
Listeners (*Global Traffic --> Listeners*)	*Internal Listeners*	
	Destination	Type the IP address on which the Global Traffic Manager listens for network traffic.
	VLAN Traffic	Select **Enabled On** from the list, and then select the *Internal* VLAN(s) and add them to the Selected list.
	Protocol	**UDP**
	DNS Profile	Select the DNS profile you created above.
	Create a second internal Listener using Protocol TCP; all other settings are the same.	
	External Listeners	
	Destination	Type the IP address on which the Global Traffic Manager listens for network traffic.
	VLAN Traffic	Select **Enabled On** from the list, and then select the *External* VLAN(s) and add them to the Selected list.
	Protocol	**UDP**
	DNS Profile	Select the DNS profile you created above.
	Create a second External Listener using Protocol TCP; all other settings are the same.	
Servers (*Global Traffic --> Servers*)	*GTM Server*	
	Name	Type a unique name.
	Address list	Type the Self IP address of this GTM system.
	Data Center	Select the Data Center where this GTM resides.
	Virtual Server Discovery	**Enabled**
	LTM Servers	
	Name	Type a unique name.
	Product	Select **BIG-IP System (Single)** or **BIG-IP System (Redundant)** as applicable.
	Address list	Type the Self IP address of this GTM system.
	Data Center	Select the Data Center where this LTM resides.
	Health Monitor	**bigip**
	Virtual Server Discovery	**Enabled**
	Repeat for each BIG-iP LTM system on which you deployed the Lync iApp template.	
	Important: After creating all of the LTM Servers, go to <u>Enabling connectivity with remote BIG-IP systems on page 9</u> and perform the commands before continuing.	

Adding to the Dependency list

For three of the pools (Edge External Reverse Proxy, Web Conferencing service and A/V service), you must add the Access Edge virtual server to the Dependency list of each member. This dependency ensures that all of the external virtual servers from the corresponding data center are marked down if no Front End servers from that data center are available to respond to SIP requests. The additional monitor you created when configuring the Lync iApp template disables the Access Edge virtual server if this health check fails, causing all other virtual servers dependent on the Access Edge virtual server to be disabled.

To add to the member Dependency List

1. On the Main tab, under **Global Traffic**, click **Servers**. The Server list opens.

2. From the list, find the row of the appropriate BIG-IP system, and click the numbered link in the Virtual Server column.

3. From the Virtual Server list, click the appropriate virtual server as described in the Dependency List section in the following table.

4. In the **Dependency List** section, select the Access Edge virtual servers from the specific Region as described in the following table.

5. Click the **Add** button.

6. Repeat Steps 4 and 5 for each virtual server you need to add to the list.

7. Click **Finished**.

8. Repeat this procedure as necessary.

The following section of the table contains BIG-IP GTM Pool and Wide IP configuration information.

Common settings for the GTM Pools
The following settings are the same across all GTM Pools in our example. Use the settings appropriate for your configuration:

- Name: Give each GTM Pool a unique name
- Load Balancing Modes: *Preferred:* **Topology**
  Alternate: **Global Availability**
  Fallback: **Return to DNS**
- Verify Virtual Server Availability: **Enabled** (this is the default, but must be enabled)
- Do **not** assign a GTM monitor to the pools.

Common settings for the GTM Wide IPs
The following setting is the same across all GTM Wide IPs:

- Load Balancing Mode: **Topology**

Lync Service	Non-default settings/Notes	
Access Service	*Pool*	
	Member List: Virtual Server	Select all LTM **Access Service** virtual servers on **port 5061**
	Wide IP	
	Name	Type the SIP domain FQDN (such as sip.example.com)
	Pool List: Pool	Select the Pool you created above
Web Conferencing Service	*Pool*	
	Member List: Virtual Server	Select all LTM **Web Conferencing Service** virtual servers on **port 443**
	Dependency List	For each Pool member, add the Access Edge virtual servers on port 5061 or 443 **from the same GTM Data Center** to the Dependency List. See *Adding to the Dependency list on page 7* for instructions.
	Wide IP	
	Name	Type the Web Conferencing Service FQDN (such as conf.example.com)
	Pool List: Pool	Select the Pool you created above

Lync Service	Non-default settings/Notes	
A/V Service	*Pool*	
	Member List: Virtual Server	Select all LTM **A/V Service** virtual servers on **port 443**
	Dependency List	For each Pool member, add the Access Edge virtual servers on port 5061 or 443 ***from the same GTM Data Center*** to the Dependency List. See *Adding to the Dependency list on page 7* for instructions.
	Wide IP	
	Name	Type the A/V Service FQDN (such as av.example.com)
	Pool List: Pool	Select the Pool you created above
XMPP Federation	*Pool*	
	Member List: Virtual Server	Select all LTM **XMPP** virtual servers on **port 5269**
	Dependency List	For each Pool member, add the Access Edge virtual servers on port 5269 ***from the same GTM Data Center*** to the Dependency List. See *Adding to the Dependency list on page 7* for instructions.
	Wide IP	
	Name	Type the XMPP Service FQDN (such as xmpp.example.com)
	Pool List: Pool	Select the Pool you created above
Meet Simple URL	*Pool (Internal Front End or Director port 5061)*	
	Member List: Virtual Server	Select all LTM **Internal Front End *or* Director** virtual servers on **port 5061**
	Pool (External Edge Reverse Proxy)	
	Member List: Virtual Server	Select all LTM **External Edge Reverse Proxy** virtual servers
	Wide IP	
	Name	Type the Meet Simple URL FQDN (such as meet.example.com)
	Pool List: Pool	Select both of the Pools you created above
Dialin Simple URL	*Pool (Internal Front End or Director port 5061)*	
	Member List: Virtual Server	Select all LTM **Internal Front End *or* Director** virtual servers on **port 5061**
	Pool (External Edge Reverse Proxy)	
	Member List: Virtual Server	Select all LTM **External Edge Reverse Proxy** virtual servers
	Wide IP	
	Name	Type the Dialin Simple URL FQDN (such as dialin.example.com)
	Pool List: Pool	Select both of the Pools you created above
External Lync Mobility	*Pool*	
	Member List: Virtual Server	Select all LTM **External Edge Reverse Proxy** virtual servers
	Wide IP	
	Name	Type the Internal Lync Mobility FQDN (such as: lyncdiscover.example.com)
	Pool List: Pool	Select the Pool you created above
Edge Pool	*Pool*	
	Member List: Virtual Server	Select all LTM **Internal Edge** virtual servers
	Wide IP	
	Name	Type the Edge Pool FQDN (such as edge.example.com)
	Pool List: Pool	Select the Pool you created above
Director (5061)	*Pool (Internal Director Web Services)*	
	Member List: Virtual Server	Select all LTM **Director** virtual servers on **port 5061**
	Pool (Reverse Proxy)	
	Member List: Virtual Server	Select all LTM **External Edge Reverse Proxy** virtual servers
	Wide IP	
	Name	Type the Director FQDN (such as dir.example.com)
	Pool List: Pool	Select both of the Pools you created above

Lync Service	Non-default settings/Notes	
Front End Web Services (5061)	*Pool (Front End port 5061)*	
	Member List: Virtual Server	Select all LTM **Internal Front End** virtual servers on **port 5061**
	Pool (Reverse Proxy port 5061)	
	Member List: Virtual Server	Select all LTM **External Edge Reverse Proxy** virtual servers.
	Dependency List	For each Pool member, add the Access Edge virtual servers on port 5061 or 443 ***from the same GTM Data Center*** to the Dependency List. See *Adding to the Dependency list on page 7* for instructions.
	Wide IP	
	Name	Type the Web Services domain FQDN (such as chat.example.com)
	Pool List: Pool	Select both of the Pools you created above
Internal Lync Mobility	*Pool*	
	Member List: Virtual Server	Select all LTM **Front End *or* Director** virtual servers on **port 443**
	Wide IP	
	Name	Type the Internal Lync Mobility FQDN (such as: lyncdiscoverinternal.example.com)
	Pool List: Pool	Select the Pool you created above

Enabling connectivity with remote BIG-IP systems

After creating the LTM Servers on the BIG-IP GTM, open a command prompt from the BIG-IP GTM, and then run the following commands for each BIG-IP LTM.

From the GTM command line, type
`big3d_install <IP address of target system>`
where the target system is the LTM that you want to add as a server on the GTM. This pushes out the newest version of big3d.

Next, type
`bigip_add`
to exchange SSL keys with the LTM. Type the password at the prompt, and then type
`iqdump <ip address of remote box>`.
If the boxes are communicating over iQuery, you see a list of configuration information from the remote BIG-IP.

The **bigip_add** command must be run for every BIG-IP in the configuration.
The following Topology Regions and Records should be configured as appropriate for your configuration. The entries in the table are examples from our configuration.

BIG-IP GTM Object	Non-default settings/Notes	
Topology Regions (*Global Traffic -->Topology -->Regions*)	*Internal*	
	Name	Type a unique name. We recommend using Internal.
	Region Members	Add Internal region members. In our example we use **IP Subnet** as the Member Type, and **is**, and then add the members of our internal subnet.
	External	
	Name	Type a unique name. We recommend using External.
	Region Members	Add External region members. In our example we use **IP Subnet** as the Member Type, and **is not**, and then add the members of our External subnet.
Topology Records (*Global Traffic -->Topology -->Records*)	*Record for internal Front End Web services requests*	
	Name	Type a unique name.
	Request Source	From the lists, select the appropriate values. In our example, we use **"Region" "is" "internal"**
	Destination	From the lists, select: **"Pool" "is"** and then select your Internal **Front End Web Services Pool.**

BIG-IP GTM Object	Non-default settings/Notes	
Topology Records (*Continued*)	*Record for internal Director Web services requests*	
	Name	Type a unique name.
	Request Source	From the lists, select the appropriate values. In our example, we use: **"Region" "is" "internal"**
	Destination	From the lists, select: **"Pool" "is"** and then select your Internal **Director Web Services Pool.**
	Record for Reverse Proxy requests	
	Name	Type a unique name.
	Request Source	From the lists, select the appropriate values. In our example, we use: **"Region" "is" "External"**
	Destination	From the lists, select: **"Pool" "is"** and then select your Internal **Reverse Proxy Pool.**
	Record for Reverse Proxy requests	
	Name	Type a unique name.
	Request Source	From the lists, select the appropriate values. In our example, we use: **"Region" "is" "internal"**
	Destination	From the lists, select: **"Pool" "is"** and then select your Internal **Front End Web Services Pool.**
	Geographical Records	
	Name	Type a unique name.
	Request Source	From the lists, select the appropriate values. In our example, we use: **"State" "is" "United States" / "New York"**
	Destination	From the lists, select the appropriate values. In our example, we use: **"Data Center" "is" "New York DC".**
	Geographical Records	
	Name	Type a unique name.
	Request Source	From the lists, select the appropriate values. In our example, we use: **"State" "is" "United States" / "New Jersey"**
	Destination	From the lists, select the appropriate values. In our example, we use: **"Data Center" "is" "New Jersey DC".**

Creating a Distributed Application for Lync

In this section, we create a Distributed Application on the BIG-IP GTM for Lync.

To create a Distributed Application, on the Main tab, expand **Global Traffic**, click **Distributed Applications**, and then click the **Create** button. Use the following table for guidance on the settings.

Setting	Non-default settings/Notes
Name	Type a unique name
Dependency Level	**None**
Persistence	Check the box to enable persistence. Configure the persistence settings as applicable for your configuration. We leave the default settings.
Member List	From the Wide IP list, select a Wide IP you created as a part of this configuration and then click the Add button. Repeat for all Wide IPs.

All Data Centers, Links, and Server objects associated with the GTM deployment for Lync are automatically added to the properties of the Distributed Application.

Manually failing over to Lync resources in another Data Center

Use the following procedure to manually fail over to Lync resources on another Data Center.

To manually fail over to Lync resources in another Data Center

1. On the Main tab, expand **Global Traffic**, and then click **Distributed Applications**.

2. On the Menu bar, click **Data Centers**.

3. Click a check in the box next to the name of the Data Center for which you would like GTM to stop sending DNS responses.

4. Click **Disable Distributed Application Traffic**.

You can also use this method to prevent traffic to a Data Center that has automatically failed over using the monitor dependencies previously created in this guide. Once the Data Center or failed Lync Servers have recovered, highlight the Data Center's Distributed Application object and click **Enable Distributed Application** to resume sending responses.

Lync Server Connection Draining

Lync Server includes a feature known as *Server Draining*. This setting prevents new connections to the server while allowing existing connections to terminate gracefully. If you are using Server Draining with Lync 2010, F5 recommends disabling the BIG-IP pool member(s) associated with that server prior to selecting the **Prevent new sessions for all services** option for that server from the Lync Control Panel. When the server is ready for traffic, reverse the process by enabling new sessions for the Lync server, and then enabling its associated pool member(s) on the BIG-IP system.

Document Revision History

Version	Description	Date
1.0	New document	05/01/2012
1.1	Added the "Why F5?" section to the first page, explaining the scenarios and advantages of using BIG-IP GTM for DNS load balancing.	05/10/2012
1.2	- Added support for Lync 2013, updated configuration example and diagrams for Lync 2013. - Added a Wide IP for the port 5269 virtual server, dependent on the Access Service Wide IP. - Added support for BIG-IP versions 11.2 - 11.4	11/05/2013