Deployment Guide

Document version: 2.6

Important: This guide has been archived. While the content in this guide is still valid for the products and versions listed in the document, it is no longer being updated and may refer to F5 or third party products or versions that have reached end-of-life or end-of-support. For a list of current guides, see https://f5.com/solutions/deployment-guides.

What's inside:

Deploying the BIG-IP LTM v10 with Microsoft Lync Server 2010 and 2013

Welcome to the Microsoft® Lync® Server 2010 deployment guide. This document contains guidance on configuring the BIG-IP® Local Traffic Manager™ (LTM) version 10.2.2 HF-1 and later versions in the v10 branch with Microsoft Lync Server 2010 and 2013.

This deployment guide is the result of collaboration and interoperability testing between Microsoft and F5 Networks using Microsoft Lync Server and the BIG-IP LTM. Organizations using the BIG-IP LTM benefit from mission-critical availability, intelligent traffic management, simple scalability, and enhanced security for Lync Server deployments.

For more information on Microsoft Lync Server see
http://www.microsoft.com/en-us/lync/default.aspx

For more information on the F5 BIG-IP LTM, see
http://www.f5.com/products/big-ip/product-modules/local-traffic-manager.html

Products and versions

Product	Version
BIG-IP LTM and Virtual Edition	10.2.2 HF-1 and later in the 10.x branch
Microsoft Lync Server	2010, 2013

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

To leave feedback for this or other F5 solution documents, email us at *solutionsfeedback@f5.com*

If you are using the BIG-IP system, version 11.0 or later, see
http://www.f5.com/pdf/deployment-guides/microsoft-lync-iapp-dg.pdf

Microsoft Partner Network

Prerequisites and configuration notes

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

- When used with Lync 2010 or 2013, a BIG-IP appliance and the BIG-IP VE (Virtual Edition) are configured in the same manner and offer the same functionality. Performance for large-scale sites will be better met with BIG-IP hardware, particularly for functions such as the Edge Web Conferencing service where SSL/TLS connections are terminated on the BIG-IP system.

- The BIG-IP system ***must be running version 10.2.2 HF-1 or a later*** version in the 10.x branch. If you are using BIG-IP LTM v11, use the v11 deployment guide found on f5.com: *http://www.f5.com/pdf/deployment-guides/microsoft-lync-iapp-dg.pdf*. BIG-IP version 10.2.2 HF1 includes a fix for an SSL handshake issue with TLS v1.1 or 1.2 that can affect Lync deployments. For more information, see *https://support.f5.com/kb/en-us/solutions/public/13000/000/sol13037.html*.

- When Microsoft documentation refers to a hardware load balancer (HLB), it is equivalent to the industry term *Application Delivery Controller (ADC)*, in this case F5's BIG-IP LTM.

- The BIG-IP LTM can be used in place of "DNS load balancing" in front of an Enterprise Edition pool of Front End servers and a pool of Director servers. Additionally, BIG-IP is supported between the Front End and Edge servers, and in front of consolidated Edge servers.

- In this guide, we assume most F5 administrators are familiar with configuring BIG-IP LTM objects, such as health monitors, pools and virtual servers, so the main section of the guide contains guidance on configuration objects, but does not contain step-by-step procedures.

  If you are unfamiliar with configuring the BIG-IP system, *Appendix A: Performing the BIG-IP configuration tasks on page 12* contains step-by-step instructions for configuring all BIG-IP objects that are a part of this guide.

- For Lync Server, do **not** use the Application Template for Microsoft Office Communications Server 2007 R2. Use the configuration described in this document, or upgrade to BIG-IP v11 and use the iApp template.

- Provision appropriate IP addresses for use in the BIG-IP virtual servers. See the Configuration tables for number of virtual servers and their Lync Server role.

- The configuration tables on the following pages specify *Automap* for SNAT configuration. With SNAT Automap configured, BIG-IP LTM translates the source IP address of each connection to that of its own self IP on the local subnet.
  As an alternative, you might want to SNAT to an address other than the self IP; for instance, you might want to be able to distinguish LTM monitor traffic (which always comes from the self IP) from application traffic. To accomplish this, you can create a *SNAT pool* containing a single, otherwise-unused IP address on the local subnet and use that in place of Automap. For more information on SNATs, see the BIG-IP documentation, available on Ask F5: *http://support.f5.com/kb/en-us/products/big-ip_ltm.html*.

- Depending on which Lync Services you are deploying, you need to know specific information from your Lync Server implementation to include when configuring the BIG-IP system. The following list shows the information you need and where to find it in the Lync Topology Builder. For more information, see the Lync documentation.
  - Define Simple URLs: **Site Properties>Simple URLs**.
  - Front End Web Services FQDNs, Hardware Load Balancer Monitoring Port, Collocated Mediation Server: **Enterprise Edition Front End Pool>Pool Properties**.
  - Director Web Services FQDNs: **Director Pools>Pool Properties**.
  - Edge Internal FQDN, Next Hop Pool, External Edge Services FQDNs and ports: **Edge Pools>Pool Properties**.
  - Specific settings for Edge: **A/V Edge service is NAT enabled**: **Not Checked**
  - Next hop selection: **Select Director pool if deploying Director Servers**
  - Enable separate FQDN and IP address for Web Conferencing and A/V: **Checked**
  - SIP Access Port: **443** or **5061**
  - Web Conferencing Edge Service Port: **443**
  - A/V Service Port: **443**

  You can run the Lync Topology Builder either before or after performing the BIG-IP configuration; however, because of the complexity of Lync deployment, F5 recommends gathering all information required by **both** the Lync Topology Builder and the BIG-IP configuration prior to beginning. For more information, see the Microsoft Lync documentation.

Configuration example

The BIG-IP LTM system can be used to add high availability and traffic direction to an Microsoft Lync Server Enterprise Pool. Additionally, the BIG-IP LTM system provides required SNAT functionality to enable inter-server communication within the pool.

The following example shows a typical configuration with a BIG-IP LTM system and an Lync Server 2010 deployment. With multiple Lync Servers in a pool there is a need for distributing the incoming session requests among the servers. Figure 1 shows a logical configuration diagram.

The example in figure 1 does not represent every possible configuration scenario. The following simplified diagrams show other possible configuration options using the BIG-IP LTM with Lync Server.

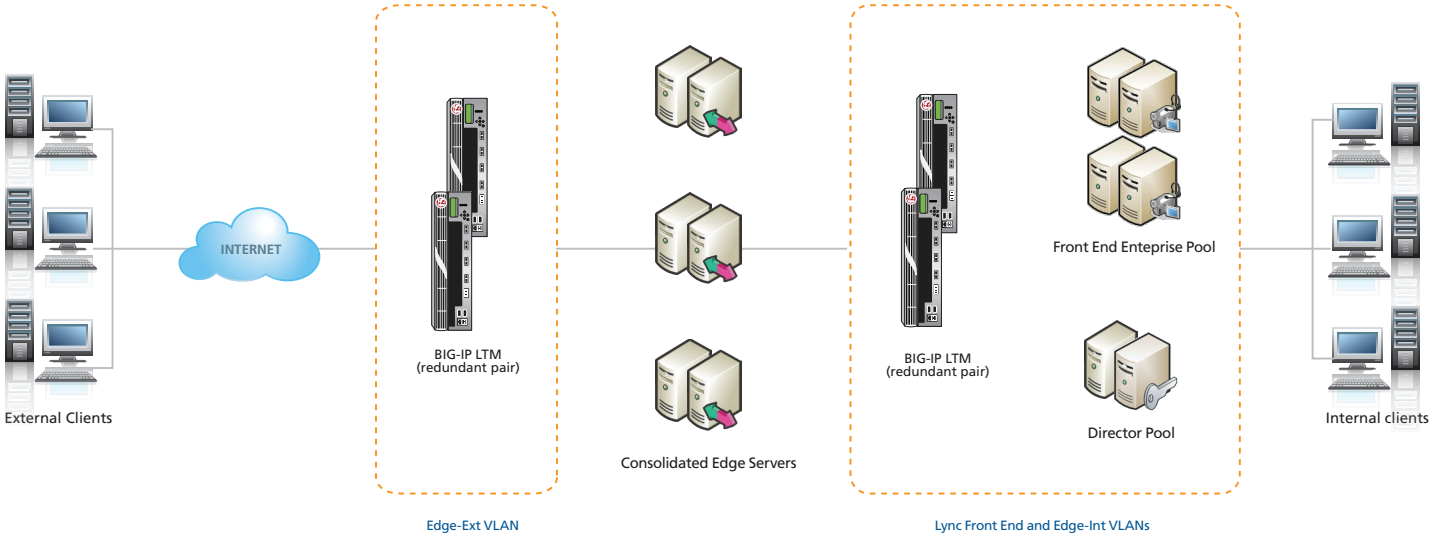

Figure 1: *Logical configuration example*

Figure 2 shows a BIG-IP LTM for External Edge Services, with a second BIG-IP LTM for Internal Edge Services, Front End Enterprise Services and Director Services. Note that while a single BIG-IP LTM is shown, we strongly recommend using LTM devices in redundant pairs.

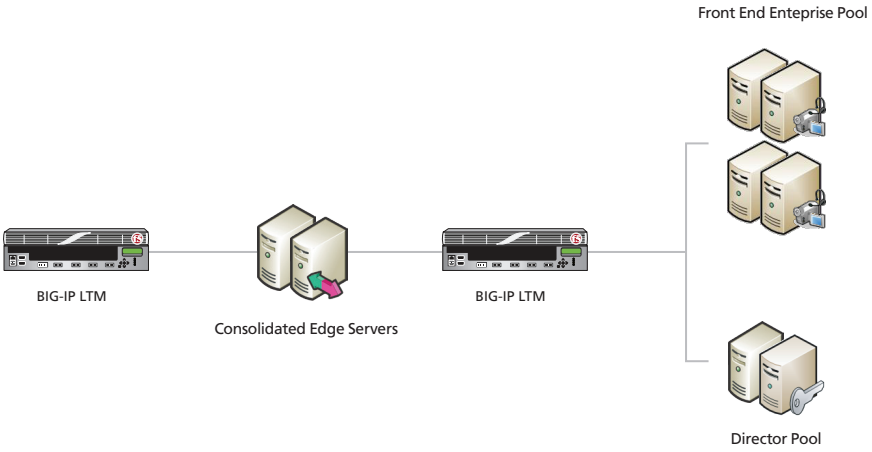

Figure 2: *Alternate logical configuration example*

Figure 3 shows a single BIG-IP LTM (redundant pair) for all internal and external Lync Server services.

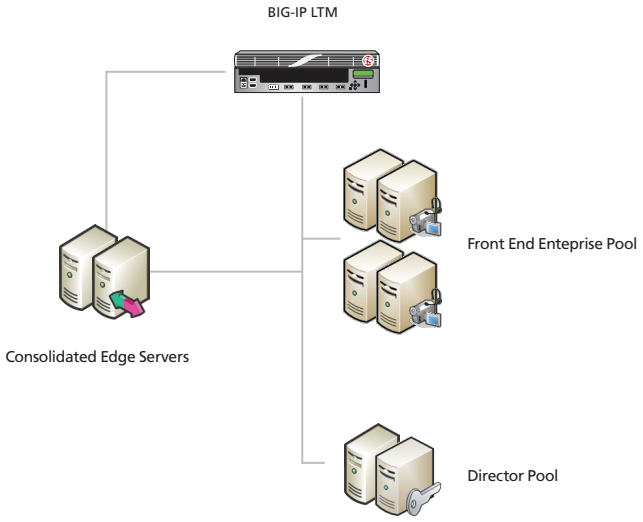

Figure 3: *Alternate logical configuration example*

Configuration Flow

This work flow diagram visually represents the BIG-IP LTM tasks listed in the configuration tables on the following pages. There is a configuration table for each Lync Role:

- *Configuration table for BIG-IP objects: Lync Front End Services on page 5*
- *Configuration table for BIG-IP objects: Lync Director Services on page 6*
- *Configuration table for BIG-IP objects: Edge Servers - External Interface on page 10*
- *Configuration table for BIG-IP objects: Edge Servers - Internal Interface on page 11*

You should follow this workflow, iterating through all rows in each of the configuration tables

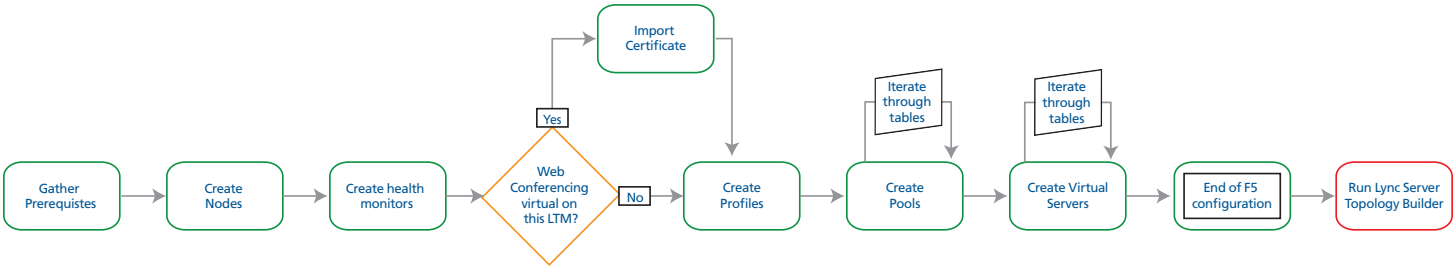

Figure 4: *Work flow diagram*

Remember, if you are unfamiliar with configuring the BIG-IP LTM, see *Appendix A: Performing the BIG-IP configuration tasks on page 12* for step-by-step instructions for configuring each type of object contained in the following tables.

Configuring the BIG-IP system for Lync Server 2010 and 2013

Use following tables to configure the BIG-IP system. This first table shows the non-default settings on BIG-IP objects for the Lync Front End Services. BIG-IP pool members (column 2) are each of the Lync Front End Server pool members (use **Least Connections (Node)** load balancing for all pools). See *Using separate internal and external BIG-IP systems versus a single BIG-IP system on page 6* for guidance on the different BIG-IP system deployment scenarios.

Configuration table for BIG-IP objects: Lync Front End Services

Virtual Server	Pool	Health monitor	Profiles	Persistence profile	SNAT enabled?	Notes
Service Port: **80**	Service Port: **80**[1] **Action on Service down: Reject**	**Lync-http-fe**: Base HTTP parent **Lync-tcp-5061-fe:**[6] Base TCP parent: - **Alias Service Port**: **5061**	***Lync-tcp-fe***: Base TCP Parent profile with **Idle Timeout** set to **1800**	**Lync-source-fe**: Source Address Affinity parent **Timeout** set to **1800**	Yes [2]	HTTP
Service Port: **135**	Service Port: **135**[1] **Action on Service down: Reject**	**lync-tcp-monitor-fe**: Base TCP parent with no required changes	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	RPC
Service Port: **443**	Service Port: **443**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe* **and** *Lync-tcp-5061-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	HTTPS
Service Port: **444**	Service Port: **444**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	
Service Port: **448**	Service Port: **448**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	
Service Port: **5061**	Service Port: **5061**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe* *Optional monitor* [3]*:* **Lync-sip-monitor-fe** Base SIP monitor - **Mode** set to **TCP**. - **Additional Accepted** - **Status Code**: add code **401** & **488** - **Alias Service Port**: **5060**		*Default*: SSL [4] **Timeout** set to **1800** *Fallback*: Source Address Affinity.	Yes [2]	SIP over TLS
Service Port: **5067**[5]	Service Port: **5067**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	This service may be collocated on your FE servers or on separate Mediation servers
Service Port: **5068**[5]	Service Port: **5068**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	Same note as above
Service Port: **5070**[5]	Service Port: **5070**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	Same note as above
Service Port: **5071**	Service Port: **5071**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	
Service Port: **5072**	Service Port: **5072**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	
Service Port: **5073**	Service Port: **5073**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	

[1] *Use the Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 17*)
[3] For the SIP monitor, additional steps need to be taken on the Microsoft Lync Front-End Servers. See *Creating a SIP monitor for the Front End servers on page 13*
[4] SSL persistence is optional but recommended
[5] These virtual servers are only necessary if deploying Lync Mediation Servers.
[6] This TCP monitor is used to support bringing down pool members when Lync servers are put into Maintenance Mode.

Virtual Server	Pool	Health monitor	Profiles	Persistence profile	SNAT enabled?	Notes
Service Port: **5075**	Service Port: **5075**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	
Service Port: **5076**	Service Port: **5076**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	
Service Port: **5080**	Service Port: **5080**[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-fe*	Use *Lync-tcp-fe*	Use *Lync-source-fe*	Yes [2]	
Service Port: **8080**	Service Port: **8080**[1] **Action on Service down: Reject**	Use *Lync-http-fe* **and** *Lync-tcp-5061-fe*[6]	Use *Lync-tcp-fe* **Client SSL**: *Lync-fe-client-ssl*: Base client SSL profile. *Important:* Must use same certificate used by Lync Server. **Server SSL**: *Lync-fe-server-ssl*: Base server SSL profile with proper certs. **HTTP**: *Lync-fe-http* Base HTTP parent with no optimizations	Use *Lync-source-fe*	Yes [2]	

Table 1: *Configuration table for Front End Server Objects*

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 17*)
[3] For the SIP monitor, additional steps need to be taken on the Microsoft Lync Front-End Servers. See *Creating a SIP monitor for the Front End servers on page 13*
[4] SSL persistence is optional but recommended
[5] These virtual servers are only necessary if deploying Lync Mediation Servers.
[6] This TCP monitor is used to support bringing down pool members when Lync servers are put into Maintenance Mode.

Configuration table for BIG-IP objects: Lync Director Services

The following table shows the non-default settings on BIG-IP LTM objects for the Lync Director services. The BIG-IP pool members (column 2) for the following table are each of the Lync Director servers.

Virtual Server port	Pool	Health monitor	TCP profiles	Persistence profile	SNAT enabled?	Notes
443	Service Port: 443[1] **Action on Service down: Reject**	**lync-tcp-monitor-dir**: Base TCP monitor with no required changes **Lync-tcp-5061-dir[3]:** Base TCP parent: **Alias Service Port**: **5061**	Standard TCP	None	Yes [2]	
444	Service Port: 444[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-dir*	Standard TCP	None	Yes [2]	
5061	Service Port: 5061[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-dir*	Standard TCP	None	Yes [2]	SIP over TLS

Table 2: *Configuration table for Director services*

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 17*)
[3] This TCP monitor is used to support bringing down pool members when Lync servers are put into Maintenance Mode.

Configuration table for BIG-IP objects when a reverse proxy is used

When deploying a Scaled Edge topology with a reverse proxy server, you need to create the following virtual servers on the BIG-IP LTM, depending on whether you are using Director servers. Additional details, including a configuration diagram, can be found at *http://technet.microsoft.com/en-us/library/gg398478.aspx*. There are internal and external BIG-IP virtual servers for the reverse proxy configuration. You can optionally create an external reverse proxy virtual server on the BIG-IP LTM that replaces the need for a separate reverse proxy device.

Internal reverse proxy configuration table

For the internal side, there are additional virtual servers between your reverse proxy and your Front End pool, or optionally your Director pool. In most cases, this is the same BIG-IP LTM you configured with the virtual servers for your Front End or Director pools.

Virtual Server port	Pool	Health monitor	Profiles	Persistence profile	SNAT enabled?	Notes
Front End reverse proxy virtual server						
4443	Front End pool members on port 4443[1] **Action on Service down: Reject**	**Lync-https-4443-fe**: Base HTTPS monitor **Alias Service Port** set to **4443** Other settings optional **Lync-tcp-5061-in-rp[3]:** Base TCP parent: **Alias Service Port**: **5061**	Use *Lync-tcp-fe* **Client SSL**: *Lync-fe-client-ssl*: Base client SSL profile. *Important:* Must use same certificate used by Lync Server. **Server SSL**: Lync-fe-server-ssl: Base server SSL profile with proper certs. **HTTP**: *Lync-fe-http* Base HTTP parent profile with no optimizations	**Important:** *Required for Lync 2010 and Lync 2013 implementations using Lync 2010 servers. It is optional for Lync 2013 only.* **Lync-cookie-fe-in-rp**: *Type*: **Cookie** *Cookie Name:* **MS-WSMAN** *Expiration:* Clear the *Session Cookie* box, and then set the Expiration to **3650 Days** **If BIG-IP v11:** *Always Send Cookie*: **Enabled** **If BIG-IP v10:** Configure the Cookie insert iRule on *page 8*	Yes [2]	**Important:** *This virtual server is only required when a reverse proxy server is deployed as part of a Lync Edge server implementation.*

Table 3: *Configuration table if using a reverse proxy*

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 17*)
[3] This TCP monitor is used to support bringing down pool members when Lync servers are put into Maintenance Mode.

This next virtual server is for the reverse proxy if you are using Director servers.

Virtual Server port	Pool	Health monitor	Profiles	Persistence profile	SNAT enabled?	Notes
Front End reverse proxy virtual server						
4443	Director server pool members on port 4443[1] **Action on Service down: Reject**	**Lync-https-4443-fe**: Base HTTPS monitor **Alias Service Port** set to **4443** Other settings optional **Lync-tcp-5061-in-rp[3]:** Base TCP parent: **Alias Service Port**: **5061**	Use *Lync-tcp-fe* **Client SSL**: *Lync-fe-client-ssl*: Base client SSL profile. *Important:* Must use same certificate used by Lync Server. **Server SSL**: Lync-fe-server-ssl: Base server SSL profile with proper certs. **HTTP**: *Lync-fe-http* Base HTTP parent profile with no optimizations	Cookie: Cookie Name set to **MS-WSMAN** **If BIG-IP v11:** *Always Send Cookie*: **Enabled** **If BIG-IP v10:** Configure the Cookie insert iRule on *page 8* **(optional for Lync 2013)**	Yes [2]	**Important:** *This virtual server is only required when a reverse proxy server is deployed as part of a Lync Edge server implementation.*

Table 3 continued: *Configuration table if using Director servers and reverse proxy*

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 17*)
[3] This TCP monitor is used to support bringing down pool members when Lync servers are put into Maintenance Mode.

NOTE: When deploying an external reverse proxy for Lync web services, F5 recommends either deploying an LTM virtual server to receive external Lync web services traffic as described in the following section, or locating the reverse proxy server (such as Microsoft Threat Management Gateway) directly on a public network. Deploying a third-party external reverse proxy server behind the BIG-IP LTM is not a supported configuration.

External reverse proxy configuration table

Create the following virtual server if you want to use the BIG-IP LTM to act as a reverse proxy and eliminate the need for a separate device. This virtual server uses an iRule to properly send traffic to the correct location.

Important: *This virtual server is only required when you want to replace a separate reverse proxy device.*

Virtual Server	Pool	Health monitor	Profiles	Persistence profile	SNAT enabled?	Other
Front End reverse proxy virtual server						
Service port: **443** **Critical:** Do **NOT** assign a default pool to this virtual server. The pool assignment is handled by the iRule.	*Front End reverse proxy pool*: The only member is the IP address of the internal Front End port 4443 virtual server you created. *Director reverse proxy pool*: If using Director servers, create an additional pool. The only member is the IP address of the internal Director port 4443 virtual server. Both use Service Port **4443**[1] **Action on Service down: Reject**	**Lync-https-4443-fe**: Base HTTPS monitor **Alias Service Port** set to **4443** Other settings optional **Lync-tcp-5061-ex-rp**[3]**:** Base TCP parent: **Alias Service Port**: **5061**	Use *Lync-tcp-fe* **Client SSL**: *Lync-fe-client-ssl*: Base client SSL profile. *Important:* Must use same certificate used by Lync Server. **Server SSL**: Lync-fe-server-ssl: Base server SSL profile with proper certs. **HTTP**: *Lync-fe-http* Base HTTP parent profile with no optimizations	None	Yes [2]	You must enable Port Translation on this virtual server (enabled by default). **Critical:** You must also attach an iRule to this virtual server. See the following section.

Table 4: *Configuration table if using Director servers and reverse proxy*

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 17*)
[3] This TCP monitor is used to support bringing down pool members when Lync servers are put into Maintenance Mode.

Creating the iRules

This section contains iRules referenced from the configuration tables. To create an iRule, from the Main tab of the BIG-IP Configuration utility, expand **Local Traffic**, click **iRules** and then click **Create**. Type a name and then copy and paste the appropriate iRule in the **Definition** section, and then click **Finished**.

Creating the Cookie Insert iRule for Lync 2010 and BIG-IP v10

The first iRule is for Front End reverse proxy virtual servers that use Cookie persistence. This iRule ensures that the cookie is sent properly.

IMPORTANT: ***This iRule is only necessary if using BIG-IP v10 AND Lync 2010***.

```
when HTTP_REQUEST {
if { [HTTP::cookie exists "MS-WSMAN"] } {
    set need_cookie 0
    } else {
    set need_cookie 1
    }
}

when HTTP_RESPONSE {
    if { ($need_cookie == 1) && (! [HTTP::cookie exists "MS-WSMAN"]) } {
    scan [IP::server_addr] "%u.%u.%u.%u" a b c d
    set e [TCP::server_port]
    set cookie "[format %u [expr {($d<<24)|($c<<16)|($b<<8)|$a}]].[expr {($e & 0xff)<<8 | ($e >>8)}].0000"
    HTTP::cookie insert name "MS-WSMAN" value $cookie path "/"
    unset cookie
    }
   unset need_cookie
}
```

Creating the iRule for the reverse proxy virtual server

For the external reverse proxy virtual server, you must create an iRule that sends traffic to the proper Lync service. The iRule you create depends on whether you are using Director servers or not, and the format of the URLs. We provide four examples in this section.

In the following examples, replace the red text with your URLs and pool names. The code goes in the Definition section when creating the iRule. The line numbers are provided for reference, do not include them in the code.

*iRule for Simple URLs in 'meet.example.com' format when you are **NOT** forwarding reverse proxy traffic to Director servers*

```
1   when HTTP_REQUEST {
2       switch -glob [string tolower [HTTP::host]] {
3            chat.example.com* { pool front_end_pool }
4            meet.example.com*  { pool front_end_pool }
5            dialin.example.com* { pool front_end_pool }
6            lyncdiscover.example.com* { pool front_end_pool }
7       }
8   }
```

*iRule for Simple URLs in 'www.example.com/meet' format when you are **NOT** forwarding reverse proxy traffic to Director servers*

```
1   when HTTP_REQUEST {
2      switch -glob [string tolower [HTTP::host]] {
3         chat.example.com* { pool front_end_pool }
4         example.com {
5         switch -glob [string tolower [HTTP::uri]] {
6            /meet* { pool front_end_pool }
7            /dialin* { pool front_end_pool }
8            }
9         }
10        lyncdiscover.example.com* { pool front_end_pool }
11     }
12  }
```

*iRule for Simple URLs in 'meet.example.com' format when you **ARE** forwarding reverse proxy traffic to Director servers*

```
1   when HTTP_REQUEST {
2       switch -glob [string tolower [HTTP::host]] {
3            chat.example.com* { pool front_end_pool }
4            dir.example.com* { pool director_pool }
5            meet.example.com*  { pool director_pool }
6            dialin.example.com* { pool director_pool }
7            lyncdiscover.example.com* { pool director_pool }
8       }
9   }
```

*iRule for Simple URLs in 'www.example.com/meet' format when you **ARE** forwarding reverse proxy traffic to Director servers*

```
1   when HTTP_REQUEST {
2       switch -glob [string tolower [HTTP::host]] {
3            chat.example.com* { pool front_end_pool }
4            dir.example.com* { pool director_pool }
5            www.example.com* {
6            switch -glob [string tolower [HTTP::uri]] {
7                 /meet* { pool director_pool }
8                 /dialin* { pool director_pool }
9                 }
10           }
11           lyncdiscover.example.com* { pool director_end_pool }
12      }
13  }
```

Attach the appropriate iRule to the virtual server.

This completes the Reverse Proxy section.

Configuration table for BIG-IP objects: Edge Servers - External Interface

The following table is for external interface of the Microsoft Lync Edge Servers. The BIG-IP pool members (column 2) are the external interface of the Lync Edge Servers

➲ **Note** *Each Lync Edge server should have a unique publicly routable IP address for each of the three Edge services (Access, A/V, and Web Conferencing) in addition to one unique public IP address for each service's BIG-IP virtual server. If you are deploying two Edge servers, you would need 9 publicly routable IP addresses.*

Virtual Server port	Pool	Health monitor	Profiles	Persistence profile	SNAT enabled?	Notes
Access Service						
Note: For the Access service, you configure either a 443 **or** a 5061 virtual server as described below. However, if you have enabled federation on port 5061 in the Lync Server Topology, and created the Access virtual server on port 443, you must **also** create the virtual server on port 5061. If you are using Lync 2013 and enabled federation with XMPP providers on port 5269 in the Lync Server Topology, you must **also** create the 5269 virtual server.						
443	Service Port: 443[1] **Action on Service down: Reject**	**lync-tcp-monitor-ext**: Base TCP monitor with no required changes	**TCP**: *Lync-edge-tcp-ext*: Base *tcp* parent profile with **Idle Timeout** set to **1800** **Nagle's Algorithm**: **Disabled**	Source Address Affinity	Yes [2]	
5061	Service Port: 5061[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-ext*	Use *Lync-edge-tcp-ext*	*Default*: SSL [3] **Timeout** set to **1800** *Fallback*: Source Address Affinity	Yes [2]	
5269	Service Port: 5269[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-ext*	Use *Lync-edge-tcp-ext*	Source Address Affinity	Yes [2]	
Web Conferencing Service						
443	Service Port: 443[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-ext*	Use *Lync-edge-tcp-ext*	Source Address Affinity	Yes [2]	
A/V Service [3]						
443	Service Port: 443[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-ext*	Use *Lync-edge-tcp-ext*	Source Address Affinity	Not recommended[4]	The A/V Edge external interfaces must have publicly routable IP addresses
3478	Service Port: 3478[1] **Action on Service down: Reject**	UDP monitor: Base UDP monitor with no required changes. ICMP monitor: Base Gateway ICMP monitor with no changes	Standard **UDP**	Source Address Affinity	Not recommended[4]	Add both monitors to the pool. The iCMP monitor ensures a pool member is properly marked down

Table 5: *Configuration table for Edge Servers - External Interface*

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method

[2] Optional, but recommended (see *Creating a SNAT on page 17*)

[3] SSL persistence is optional but recommended

[4] For best performance, F5 does not recommend SNAT for Edge A/V services. However, SNAT for these services is supported in deployments where it is required.

Configuration table for BIG-IP objects: Edge Servers - Internal Interface

The following table is for internal interface of the Microsoft Lync Edge Servers.
The BIG-IP pool members (column 2) for the following table are the internal interface of the Lync Edge Servers.

Virtual Server Port	Pool	Health monitor	Profiles	Persistence Profile	SNAT enabled?	Notes
443	Service Port: 443[1] **Action on Service down: Reject**	**lync-tcp-monitor-int**: Base TCP monitor with no required changes	***TCP:*** *Lync-edge-tcp-int*: Base *tcp* Parent profile with **Idle Timeout** set to **1800**	Source Address Affinity	Yes [2]	
3478	Service Port: 3478[1] (UDP) **Action on Service down: Reject**	UDP monitor: Base UDP monitor with no required changes	Standard **UDP**	Source Address Affinity	Yes [2]	STUN/UDP inbound/ outbound
5061	Service Port: 5061[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-int*	Use *Lync-edge-tcp-int*	*Default*: SSL [3] **Timeout** set to **1800** *Fallback*: Source Address Affinity	Yes [2]	
5062	Service Port: 5062[1] **Action on Service down: Reject**	Use *Lync-tcp-monitor-int*	Use *Lync-edge-tcp-int*	*Default*: SSL [3] **Timeout** set to **1800** *Fallback*: Source Address Affinity	Yes [2]	

Table 6: *Configuration table for Edge Servers - Internal Interface*

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 17*)
[3] SSL persistence is optional but recommended

Appendix A: Performing the BIG-IP configuration tasks

In this section, we provide step-by-step configuration procedures for those who are unfamiliar with creating objects on the BIG-IP system. You need the information from the Configuration tables to use when configuring most of the following objects.

Performing the initial configuration tasks

The initial configuration includes VLAN and Self IP address configuration. For more detailed information on initial configuration tasks on the BIG-IP LTM system, see the ***TMOS Management Guide for BIG-IP Systems*** available on Ask F5 (*http://support.f5.com/kb/en-us.html*).

Creating the VLANs

The first procedure in this deployment is to create a VLAN on the BIG-IP LTM system. Depending on the desired network architecture, you may have one or multiple VLANs associated with the BIG-IP LTM configuration:

- *One armed configuration*
  When the Lync clients reside on the same network as the Lync Front End servers, or you wish to have your BIG-IP LTM virtual servers reside on the same network as your Front End servers, you only need one VLAN. This is also known as a one armed configuration.

➲ **Note** In deployments with more than 65,000 simultaneous connections, you need to configure more than one SNAT address on the BIG-IP LTM. See *Creating a SNAT on page 17*.

- *Routed configuration*
  A more common example is when the Lync Server clients and the IP addresses of your BIG-IP LTM virtual servers reside on a different network than the Lync Front End servers. In this case, you will need an external VLAN for the incoming clients, and an internal VLAN for the Front End servers. This is known as a routed configuration.

To create a VLAN

1. On the Main tab, expand **Network**, and then click **VLANs**.
2. Click the **Create** button.
3. In the **Name** box, type a unique name for the VLAN. For example, **Edge-External-vlan**.
4. In the **Tag** box, we assign a tag. In our example, we find that using VLAN tags make management easier. However, tagging is not mandatory if your configuration can support individual interfaces instead of VLANs.
5. In the Resources section, from the **Available** list, select the interface that will have access to tagged traffic, and add it to the Tagged box by clicking the Add (>>) button.
6. Click the **Finished** button.
7. If you are using a routed configuration, you need at least two VLANs on the BIG-IP LTM system. Use the preceding procedure to create each VLAN. Give each VLANs a distinct name (such as lync-internal-vlan and lync-external-vlan), and assign them to the interfaces through which each VLANs traffic should flow.

Creating self IP addresses on the BIG-IP

Self IP addresses are the IP addresses owned by the BIG-IP LTM system that you use to access the VLANs. The next task in this configuration is to create the self IP addresses.

To create the self IP addresses

1. On the Main tab, expand **Network**, and then click **Self IPs**.
2. Click the **Create** button.
3. In the **IP Address** box, type a static IP address that resides on the VLAN you created in the preceding procedure. For example **192.168.64.245**.

4. In the **Netmask** box, type the corresponding subnet mask. For example, **255.255.255.0**.
5. From the **VLAN** list, select the appropriate VLAN you created in *Creating the VLANs on page 12Creating the VLANs on page 12*. For example **Edge-External-vlan**.
6. Click **Finished**.
7. In a routed configuration, repeat this procedure to add a self IP to each VLAN you created.

This concludes the initial configuration tasks.

Creating the application objects on BIG-IP LTM for Lync server

The next task is to create the application specific objects on the BIG-IP system for Lync server. You must repeat many of these procedures multiple times, according to the configuration tables starting with the *Configuration table for BIG-IP objects: Lync Front End Services on page 5*.

Creating the health monitors

The first task is to create the health monitors. To create the health monitor, use the following procedure and the value from the Health Monitor column of the tables above.

To configure a health monitor

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**.
2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a unique name for this monitor. We recommend using a name that includes the service or service port.
4. From the **Type** list, select the appropriate monitor type. For most Lync services, you select **TCP.**
5. All other configuration settings are optional, configure as applicable for your deployment. We recommend at least a 1:3 +1 ratio between the **Interval** and the **Timeout** if you change those values. For example, **30** and **91**.
6. Click the **Repeat** button.
7. Repeat this procedure for each monitor in the tables above.

Creating a SIP monitor for the Front End servers

By default, SIP traffic on Front End servers is encrypted on port 5061. You may optionally enable unencrypted port 5060 for the purposes of health monitoring only; normal SIP communication cannot occur on the unencrypted port. A SIP monitor is more accurate than a simple TCP monitor, which only determines whether a port is active and not if the associated service is actually running.

In addition to configuring the SIP monitor on the BIG-IP LTM, you must also modify the Lync Front End Server configuration to enable for 5060.

To enable port 5060, use the Lync Server 2010 Topology Builder to modify the properties for your Enterprise Edition Front End Pool. Select **Enable Hardware Load Balancer monitoring port** as shown in the following figure, and then choose the default port number of **5060** or enter a custom port. Port 5060 is standard for SIP; if you select another port number, it must be one that is not otherwise in use on your Front End servers, you must make sure it is permitted on the local firewalls of those servers, and you must adjust the BIG-IP LTM monitor. Re-run the Lync Server Deployment Wizard on each Front End server to apply the change.

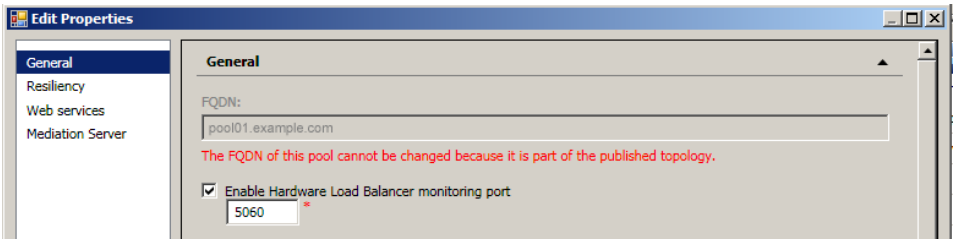

Figure 5: *Editing the General properties*

To create the BIG-IP LTM SIP monitor

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**.
2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a unique name for this monitor. We type **Lync-sip-monitor-fe**.
4. From the **Type** list, select **SIP**.
5. From the **Configuration** list, select **Advanced**.
6. From the **Mode** list, select **TCP**.
7. From the **Additional Accepted Status Codes** list, select **Status Code List**, and then type **488** in the **Status code** box. Click **Add**.
8. In the **Alias Service Port** box, type **5060** (or the custom port you selected in the Topology Builder).
9. Click **Finished**.

➲ ***Additional Information*:**
When a Hardware Load Balancer monitoring port is configured using Topology Builder, Lync Server 2010 will respond to SIP requests on that port with a status code of "488" and a reason "Port is configured for health monitoring only". The BIG-IP LTM health monitor you configured in this step treats that as an expected response from the Front End SIP service and marks the pool member as available to accept traffic.

Creating the nodes

The next task in this configuration is to create the nodes. Because so many Lync services run on the same device/node, it is more efficient to pre-configure the nodes, and then simply select the proper nodes from the node list when configuring the pools.

To create the nodes

1. On the Main tab, expand **Local Traffic**, and then click **Nodes**.
2. Click the **Create** button.
3. In the **Address** box, type the IP address of the first Lync Server node.
4. In the **Name** box, type a descriptive name for this node. We recommend including the Lync Server type (for example *front-end* or *edge)*. For example **lync-front-end-node1**.
5. Leave the values in the Configuration section at the default levels.
6. Click the **Repeat** button
7. Repeat steps 3-6 for each Lync Server node in your configuration.
8. Click the **Finished** button after completing all nodes.

Creating the pools

The next task is to create the pools. You create a pool for each of the Lync services in the tables above, using the value from the **Pool member port** column.

To create the pools

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
2. Click the **Create** button.
3. From the **Configuration** list, select **Advanced**.
4. In the **Name** box, type a unique name for this Pool. Again, we recommend using a name that includes the service or service port.

5. In the *Health Monitors* section, from the **Available** list, select the name of the monitor you created in *Creating the health monitors on page 13*, and click the Add (**<<**) button.
6. From the **Action on Service Down** list, select **Reject**.
7. *Optional*: In the **Slow Ramp Time** box, type **300**.
   Because we are using the Least Connections load balancing method, we set the Slow Ramp Time in order to ensure that if a pool member becomes available after maintenance or a new member is added, the BIG-IP does not send all new connections to that member (as a newly available member always has the least number of connections).
8. In the Resources section, from the **Load Balancing Method** list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network). For this configuration, we recommend selecting **Least Connections (member)**. In Least Connections mode, the BIG-IP LTM system passes a new connection to the node that has the least number of current connections. Least Connections mode works best in environments where the servers or other equipment you are load balancing have similar capabilities. With Lync Server, traffic from servers to clients is roughly the same on each connection.
9. In the New Members section, you add the Lync servers to the pool.
   a. Click the **Node list** option button.
   b. From the **Address** list, select the appropriate node you created in the preceding procedure from the list.
   c. In the **Service Port** box, type the service port number for this device. Refer to the appropriate configuration table for the pool you are creating, such as the *Configuration table for BIG-IP objects: Lync Front End Services on page 5*.
   d. Click the **Add** button to add the member to the list.
   e. Repeat steps b-d for each device you want to add to the pool.
10. Click the **Repeat** button. Repeat this procedure for each BIG-IP pool you need to create, based on the appropriate configuration table, such as the *Configuration table for BIG-IP objects: Lync Front End Services on page 5*.

Creating the TCP profile

In this section, we create a TCP profile. Be sure to check the TCP profile column of the tables above for specific settings in steps 5 and 6.

To create the TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, from the **Protocol** menu, select **TCP**.
3. Click the **Create** button.
4. In the **Name** box, type a name for this profile. For example, we type **Lync-tcp-fe**.
5. *For Front End servers and select other services only*: From the **Idle Timeout** row, click the **Custom** box, and then type **1800** in the Idle Timeout box.
   Important: **Steps 5 and 6 are not for all TCP profiles; be sure to check the Profiles columns of the tables above.**
6. *Only for each of the Edge Server - External Interface TCP profiles except A/V on port 3478*: From the **Nagle's Algorithm** row, click the **Custom** box, and then select **Disabled** from the list.
7. Leave the other settings at the default levels.
8. Click the **Finished** button.
9. Repeat this procedure as applicable.

Creating the UDP profile

In this section, we create a UDP profile for use with select Lync services. Again, refer to the tables above.

To create the UDP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, from the **Protocol** menu, select **UDP**.
3. Click the **Create** button.
4. In the **Name** box, type a name for this profile. In our example, we type **Lync-UDP**.
5. Configure any other options as applicable. In our example, we leave the defaults.
6. Click the **Finished** button.
7. Repeat this procedure as applicable.

Creating the Source Address Affinity persistence profile

The next profile we create is the persistence profile. For most Lync services use the Source Address Affinity (source_addr) persistence profile. Be sure to check the Persistence profile column of the tables above.

To create the persistence profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, click **Persistence**.
3. Click the **Create** button.
4. In the **Name** box, type a name.
5. From the **Persistence Type** list, select **Source Address Affinity**.
6. Configure any of the options as applicable. In our example, we leave the defaults.
7. Click **Finished**.
8. Repeat this procedure as applicable.

Creating the Optional SSL persistence profile

The SSL persistence profile is an optional profile that uses the SSL session ID for persistence. This profile should only be used with the Front End and Edge virtual servers running on port 5061.

To create the persistence profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, click **Persistence**.
3. Click the **Create** button.
4. In the **Name** box, type a name. For example, **edge-ssl-persist**.
5. From the **Persistence Type** list, select **SSL**.
6. From the **Timeout** row, click the **Custom** box, and then type **1800**.
7. Configure any of the rest of the options as applicable. We leave the defaults.
8. Click **Finished**.
9. Repeat this procedure as applicable.

Creating a Client SSL profile

The next step in this configuration is to create a Client SSL profile. This profile contains the SSL certificate and Key information for decrypting the SSL traffic on behalf of the servers. This is only necessary on select Lync services, check the tables above. You should already have the

appropriate certificate and key pairs on the BIG-IP LTM. For information on importing certificates, see the BIG-IP documentation or online help.

To create a new Client SSL profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, from the **SSL** menu, select **Client**.
3. Click the **Create** button.
4. In the **Name** box, type a name for this profile. For example, **edge-access-ssl**.
5. In the Configuration section, check the **Certificate** and **Key** Custom boxes.
6. From the **Certificate** list, select the name of the Certificate you imported for this profile.
7. From the **Key** list, select the key you imported for this profile.
8. Click the **Finished** button.
9. Repeat this procedure as applicable.

Creating a SNAT

A source network address translation (SNAT) allows for inter-server communication and provides the ability to perform certain Lync Server pool-level management operations from the servers in a pool. Additionally, in a one-armed configuration, a SNAT allows virtual servers to exist on the same IP subnet as the Lync Server hosts.

A default SNAT is appropriate for most deployments. If more than 65,000 simultaneous users are connecting to the Lync Server deployment, see "*Configuring a SNAT for large Lync Server deployments*".

Use the procedure most applicable for your deployment.

As mentioned in the prerequisites, we typically recommend *Automap* for SNAT configuration. With SNAT Automap configured, BIG-IP LTM translates the source IP address of each connection to that of its own self IP on the local subnet. As an alternative, you might want to SNAT to an address other than the self IP; for instance, you might want to be able to distinguish LTM monitor traffic (which always comes from the self IP) from application traffic. To accomplish this, you can create a *SNAT pool* containing a single, otherwise-unused IP address on the local subnet and use that in place of Automap (see Creating a SNAT pool on the following page). For more information on SNATs, see the BIG-IP LTM documentation, available on Ask F5:
http://support.f5.com/kb/en-us/products/big-ip_ltm.html.

Creating a default SNAT for less than 64,000 concurrent users

Use this procedure if your Lync Server deployment has fewer than 64,000 simultaneous users.

To create a default SNAT

1. On the Main tab, expand **Local Traffic**, and then click **SNATs**.
2. Click the **Create** button.
3. In the **Name** box, type a name. In our example, we type **lync-default-snat**.
4. From the **Translation** list, select a setting appropriate for your configuration. In our example, we select **Automap**.
5. From the **VLAN Traffic** list, select **Enabled on**.
6. In the **VLAN List** row, from the **Available** list, select the VLANs on which your Lync Servers reside, and then click the Add (**<<**) button.
7. Click the **Finished** button.

Configuring a SNAT for large Lync Server deployments

For large deployments (with 65,000 simultaneous users), we create a SNAT pool. A SNAT pool is a pool with one unused IP address, on the same subnet as the virtual servers and Lync Servers. You must create a SNAT pool for each 65,000 clients (or fraction thereof).

➲ **Important** *This procedure is only necessary for large deployments. If your Lync deployment has less than 65,000 simultaneous users, you do not need to create a SNAT pool. Use the previous procedure.*

To create a SNAT pool for large deployments

1. On the Main tab, expand **Local Traffic**, and then click **SNATs**.
2. On the Menu bar, click **SNAT Pool List**.
3. Click the **Create** button.
4. In the **Name** box, type a name for this SNAT Pool. In our example, we type **lync-snat-pool**.
5. In the **IP Address** box, type in a valid, otherwise-unused address on the subnet containing your Front End servers, and then click **Add**. Repeat this step for each additional address needed. At least one address should be added for each 65,000 anticipated concurrent connections (the number of connection generally corresponds to the number of clients).
6. Click the **Finished** button.

The next part of the SNAT pool configuration is to configure a default SNAT that uses the SNAT pool.

7. On the Main tab, expand **Local Traffic**, and then click **SNATs**.
8. Click the **Create** button.
9. In the **Name** box, type a name for this SNAT. In our example, we type **lync-default-snat**.
10. From the **Translation** list, select **SNAT Pool**.
11. From the **Select** list, select the name of the SNAT pool you created in the preceding procedure. In our example, we select **lync-snat-pool**.
12. From the **VLAN Traffic** list, select **Enabled on**.
13. In the VLAN List row, from the **Available** list, select the VLANs on which your Lync devices reside, and click the Add (<<) button.
14. Click the **Finished** button.

Creating the virtual servers

A virtual server with its virtual address is the visible, routable entity through which the Lync Servers in a load balancing pool are made available to the client (the IP address to give clients or add to DNS). In the Lync Server documentation, this is referred to as the Hardware Load Balancer (HLB) virtual IP (VIP).

The next task is to define a virtual server that references the profile and pool you created.

To create the virtual servers

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this virtual server. Again, we recommend using a name that includes the service or service port.
4. In the Destination section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server.
6. In the **Service Port** box, type the appropriate service port for this virtual server. Refer to the appropriate configuration table for the virtual server you are creating, such as the *Configuration table for BIG-IP objects: Lync Front End Services on page 5*.

7. From the **Configuration** list, select **Advanced**.
8. *For the Edge Server virtual servers on Port 3478* **only** - From the **Protocol** list, select **UDP**.
9. From the **Protocol Profile (Client)** list, select TCP profile you created in *Creating the TCP profile on page 15*.
   For the Edge Server virtual servers on Port 3478 only, select the profile you created in *Creating the UDP profile on page 15.*
10. If applicable: From the **SSL Profile (Client)** list, select the profile you created in *Creating a Client SSL profile on page 16.*
11. From the **SNAT Pool** list, select **Automap**. Remember, SNAT is required for one-armed configuration and optional for routed configuration.

12. In the Resources section, from the **Default Pool** list, select the pool you created in *Creating the nodes on page 14*.
13. From the **Default Persistence Profile** list, select the profile you created in *Creating the Source Address Affinity persistence profile on page 16* or *Creating the Optional SSL persistence profile on page 16.*
14. *If you used an SSL persistence profile as the default for select virtual servers:* From the **Fallback Persistence Profile** list, select the profile you created in *Creating the Source Address Affinity persistence profile on page 16*.
15. Click the **Finished** button.
16. Repeat this procedure for each virtual server.

This completes the BIG-IP LTM configuration.

Revision History

Version	Description	Date
1.0	New version based on the Release Candidate	N/A
1.1	Added support for the RTM version of Lync Server 2010	N/A
1.2	Added an HTTP profile to the services using Cookie Persistence	N/A
1.3	- Removed the port 8057 virtual server from the *Edge Server - Internal Interface* table. - Moved the port 4443 virtual server from the *Edge Server - Internal Interface* table to its own table on *page 18* and clarified configuration scenario. - Added SSL offload and cookie persistence guidance for the port 8080 Front End virtual server	N/A
1.4	- Removed unnecessary Front End virtual servers on ports 80, 5060, and 5069. - Added notes on configuring an optional SIP monitor. - Modified Edge Server virtual servers to have port 443 pool members. - Corrected port 3478 monitor type. - Added port 5062 virtual server. - Moved Revision History to the end of document.	N/A
1.5	Changed the guidance for SNAT on the Edge Server External Interface from No to Yes.	N/A
1.6	Changed the guidance for SNAT as follows: - SNAT is required on all Front-End and Director services. - SNAT is required for the virtuals that direct traffic to the Internal interfaces of the Edge servers. - SNAT is optional but recommended for the External interface Edge Access and Web Conferencing services. - You must NOT use SNAT on the Edge A/V Service.	N/A
1.7	- Modified the Edge Server - External Interface Access Service virtual server on port 443 to use a TCP monitor and remove the HTTP profile. The previous monitor would properly report up/down status, but the TCP monitor more accurately monitors the service. The HTTP profile was extraneous. - Added support for BIG-IP LTM version 10.2.2	N/A
1.8	- Modified the Edge Server - External Interface Access Service virtual server on port 443 to remove the Client and Server SSL profiles. Additionally changed the Persistence profile type from Cookie to Source Address Affinity. - Added configuration for enabling automatic topology replication to the Edge Server Internal Interface section. - Added a OneConnect profile to the following virtual servers: Front End 8080, Reverse proxy 4443, and Web Conferencing 443. - Added a virtual server for Edge Server Internal Interface on port 8057. - Added load balancing method recommendation to each of the tables	N/A
1.9	Added an important note stating you must be running BIG-IP version 10.2.2 HF-1 or later. Removed support for previous versions. BIG-IP version 10.2.2 HF1 includes a fix for an SSL handshake issue with TLS v1.1 or 1.2 that can affect Lync deployments. See *https://support.f5.com/kb/en-us/solutions/public/13000/000/sol13037.html* for more information	03-12-2012
2.0	- Significantly expanded *Configuration table for BIG-IP objects when a reverse proxy is used on page 7,* to include both external and internal reverse proxy configuration. - Removed the Alias Service Port from the Front End Monitor on port 8080. - Modified the Web Conferencing service health monitor from HTTPS to TCP. - Added a note on the Access Service if Federation has been enabled.	08-08-2012
2.1	Added a port 80 virtual server to the Lync Front End services table.	08-28-2012
2.2	In the Reverse proxy configuration table on page 7, removed the Always Send Cookie: Enabled setting, as this is not an option in the 10.x code branch. No further modifications to the configuration are necessary.	10-17-2012
2.3	- Added Status Code 401 to the Front End Server health monitor for port 5061, on page 5. - Added an iRule to the configuration tables for Front End and Director servers on ports 4443 and 8080 that use Cookie persistence. This iRule ensures that the cookie is sent properly. See *Creating the Cookie Insert iRule for Lync 2010 and BIG-IP v10 on page 8.*	02-04-2013

Version	Description	Date
2.4	- Added a TCP monitor on port 5061 to the HTTP pools (for the virtual servers on ports 80, 8080, 443, and 4443) to support bringing down members when Lync servers are put into Maintenance mode. - Added a timeout of 1800 to the Source Address persistence profile for the Front End servers in all cases. - Modified the port 80 and 443 Front End HTTP virtual servers to use Source Address persistence, and not cookie persistence. - Removed the port 5060 virtual server for the Director Servers as it was not used - Removed all OneConnect and NTLM profiles from the configuration.	02-15-2013
2.5	- Added Action on Service Down to Reject for all load balancing pools. - Added a note to the Reverse Proxy section on page 8 stating that deploying a third-party external reverse proxy server behind the BIG-IP LTM is not supported. - Added a timeout value of 3650 days to the cookie persistence profiles in the *Configuration table for BIG-IP objects when a reverse proxy is used on page 7*. - Modified the www.example.com/meet iRules in *Creating the iRule for the reverse proxy virtual server on page 9* to include a wildcard (*) after /meet and /dialin.	03-14-2013
2.6	- Added support for Lync Server 2013. - Updated the configuration tables for 2013. Major changes are that persistence profiles are no longer required but optional for the Reverse Proxy configuration, and added an additional virtual server in the Edge Server External Interface section for XMPP federation objects.	02-06-2014

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