Deploying the BIG-IP LTM with Microsoft Skype for Business

Welcome to the Microsoft® Skype for Business Server deployment guide. This document contains guidance on configuring the BIG-IP® Local Traffic Manager™ (LTM) with Skype for Business Server (formerly Microsoft Lync Server). BIG-IP version 11.0 introduced iApp™ Application templates, an extremely easy way to configure the BIG-IP system for Microsoft Skype for Business Server.

Why F5?

This deployment guide is the result of collaboration and interoperability testing between Microsoft and F5 Networks using Microsoft Skype for Business Server and the BIG-IP LTM. Microsoft ***requires*** hardware load balancing for Skype for Business Web Services. Organizations using the BIG-IP LTM benefit from mission-critical availability, intelligent traffic management, simple scalability, and enhanced security for Skype for Business Server deployments. This deployment also describes how to use the BIG-IP system as a reverse proxy, eliminating the need for a separate reverse proxy device.

For more information on Microsoft Skype for Business Server see *http://www.skype.com/en/business/skype-for-business/*

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 tested

Product	Version
BIG-IP LTM and Virtual Edition	Manual configuration: v11.0 - 11.6 iApp Template: 11.2 - 11.6
Microsoft Skype for Business Server	2015 RTM version
iApp Template Version	f5.microsoft_skype_server_2015.v1.0.0rc1
Deployment Guide version	1.0 (see see *Revision History on page 46*)
Last updated	07-06-2015

➲ **Important:** *Make sure you are using the most recent version of this deployment guide: http://www.f5.com/pdf/deployment-guides/microsoft-skype-for-business-iapp-dg.pdf*

Microsoft Partner Network

Contents

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 Microsoft Skype for Business acts as the single-point interface for building, managing, and monitoring Microsoft Skype for Business Server. For more information on iApp, see the *F5 iApp: Moving Application Delivery Beyond the Network* White Paper: *http://www.f5.com/pdf/white-papers/f5-iapp-wp.pdf*.

Prerequisites and configuration notes

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

BIG-IP system and general prerequisites

- **Critical:** Do not use f5.microsoft_lync_server_2010 version of the iApp template that ships with the BIG-IP system by default. F5 has released an updated version of the iApp template for Microsoft Skype for Business, which must be downloaded as shown in this document. This guide is based the new iApp.
- For users familiar with the BIG-IP system, there is manual configuration guidance at the end of this guide. However, because of the complexity of this configuration, we recommend most users deploy using the iApp template.
- Skype for Business Server uses Microsoft Office Web Apps Server 2013 for sharing Microsoft PowerPoint presentations between computers running the Skype client. F5 provides detailed configuration steps for deploying, securing, and optimizing Office Web Apps 2013 in the deployment guide available on f5.com: *http://www.f5.com/pdf/deployment-guides/microsoft-office-web-apps-dg.pdf*.
- When used with Skype for Business Server, 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 is 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 LTM.
- Microsoft documentation refers to a hardware load balancer (HLB), this is the equivalent to the industry term *Application Delivery Controller (ADC)*, in this case F5's BIG-IP LTM.
- **Critical:** *If you are using a BIG-IP version prior to 11.2 only:*
  Because this iApp uses an HTTP and SIP monitor, if you disable Strict Updates after completing the iApp configuration, the monitors stop sending requests, and mark the nodes down. At this time, we do not recommend disabling Strict Updates. If you find you need a part of the configuration that is not present in the iApp, use *Appendix: Manual Configuration table for BIG-IP objects on page 31.*
- 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. Also, BIG-IP LTM is supported between the Front End and Edge servers, and in front of Edge servers.
- You must provision appropriate IP addresses for use in the BIG-IP virtual servers. See the Configuration tables for number of virtual servers and their Skype for Business Server role.
- For an in-depth look at load balancing Skype for Business Edge servers, see *https://devcentral.f5.com/articles/the-hopefully-definitive-guide-to-load-balancing-lync-edge-servers-with-a-hardware-load-balancer*
- Deploying a third-party external reverse proxy server behind the BIG-IP LTM is not a supported configuration.

Skype for Business Server prerequisites

- Depending on which Skype for Business Services you are deploying, you need to know specific information from your Skype for Business 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 Topology Builder. For more information, see the Skype 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**

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

- If you have Skype for Business clients who will be connecting through the Edge external A/V UDP virtual server, be sure to see *Troubleshooting on page 28.*

Configuration examples

The BIG-IP LTM system can be used to add high availability and traffic direction to an Skype for Business 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 a Skype for Business Server deployment. With multiple Skype for Business Servers in a pool there is a need for distributing the incoming session requests among the servers. Figure 1 shows a logical configuration diagram.

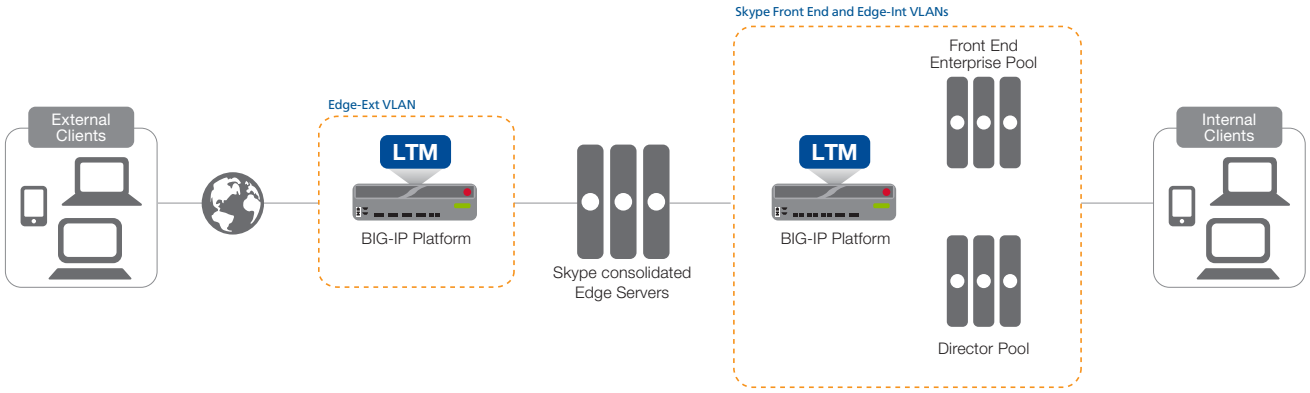

Figure 1: *Logical configuration example*

The following simplified diagram show another possible configuration option using the BIG-IP LTM with Skype for Business Server available in the iApp template.

Figure 2 shows a single BIG-IP LTM (redundant pair) for all internal and external Skype for Business Server services.

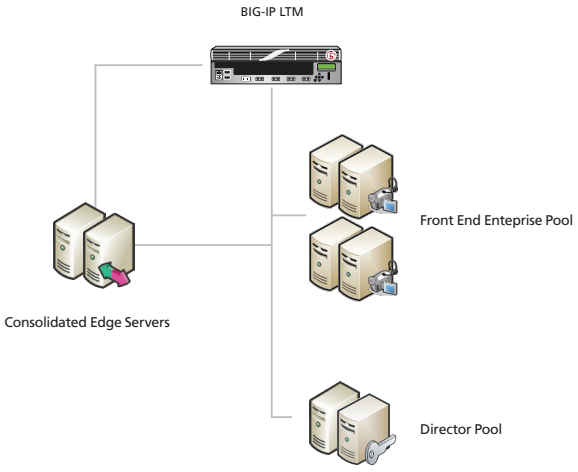

Figure 2: *Alternate logical configuration example*

Using this guide

This deployment guide is intended to help users deploy Microsoft Skype for Business Server using the BIG-IP system. This document contains guidance configuring the BIG-IP system using the iApp template, as well as manually configuring the BIG-IP system.

Using this document for guidance on configuring the iApp template

We recommend using the iApp template to configure the BIG-IP system for your Skype for Business implementation. The majority of this guide describes the iApp template and the different options the template provides for configuring the system for Skype for Business Server.

The iApp template configuration portion of this guide walks you through the entire iApp, giving detailed information not found in the iApp or online help. The questions in the UI for the iApp template itself are all displayed in a table and at the same level. In this guide, we have grouped related questions and answers in a series of lists. Questions are part of an ordered list and are underlined and in italics or bold italics. Options or answers are part of a bulleted list, and in bold. Questions with dependencies on other questions are shown nested under the top level question, as shown in the following example:

1. ***Top-level question found in the iApp template***
    - *Select an object you already created from the list* (such as a profile or pool; not present on all questions. Shown in bold italic)
    - **Choice #1** (in a drop-down list)
    - **Choice #2** (in the list)
        a. *Second level question dependent on selecting choice #2*
            - **Sub choice #1**
            - **Sub choice #2**
                a. *Third level question dependent on sub choice #2*
                    - **Sub-sub choice**
                    - **Sub-sub #2**
                        a. *Fourth level question (rare)*

Manually configuring the BIG-IP system

Users already familiar with the BIG-IP system can use the manual configuration tables to configure the BIG-IP system for the Skype for Business implementation. These configuration tables only show the configuration objects and any non-default settings recommended by F5, and do not contain procedures on specifically how to configure those options in the Configuration utility. See *Appendix: Manual Configuration table for BIG-IP objects on page 31.*

Configuring the iApp for Microsoft Skype for Business Server

Use the following guidance to help you configure Microsoft Skype for Business Server using the BIG-IP iApp template. You must have downloaded and imported the new Skype for Business iApp before beginning. See *Downloading and importing the Skype for Business iApp on page 7.*

Using separate internal and external BIG-IP systems versus a single BIG-IP system

You can use the iApp template to configure BIG-IP devices whether your Skype for Business implementation is using a single BIG-IP system (or redundant pair), or separate internal and external BIG-IP systems. The following sections provide guidance on which sections you need to configure in the iApp on which BIG-IP systems.

Using separate internal and external BIG-IP systems

If you are deploying Skype for Business Server on multiple standalone or redundant BIG-IP systems, as shown in the logical configuration example on page 3, you must complete these sections of the iApp template on each respective BIG-IP system:

- *DMZ/Perimeter Network BIG-IP system* (if deploying Skype for Business Edge services):
  - Microsoft Skype for Business Server Edge Virtual Servers - External Interface
  - Edge Server Pools - External Interface
  - Microsoft Skype for Business Server Reverse Proxy - Reverse Proxy > Forward Reverse Proxy client traffic to another BIG-IP system (if using the BIG-IP system as an external reverse proxy)
- *Internal Network BIG-IP system*
  - Microsoft Skype for Business Server Front End Virtual Servers
  - Front End Server Pools
  - Microsoft Skype for Business Server Director Virtual Servers (if deploying Director servers)
  - Director server Pools (if deploying Director servers)
  - Microsoft Skype for Business Server Edge Virtual Servers - Internal Interface (if deploying Edge services)
  - Edge Server Pools - Internal Interface (if deploying Edge services)
  - Microsoft Skype for Business Server Reverse Proxy - Reverse Proxy > Receive Reverse Proxy traffic from another BIG-IP system (if deploying Edge services and if using BIG-IP to receive Skype for Business reverse proxy traffic from another BIG-IP system or third-party reverse proxy server)

Using a single BIG-IP system (or redundant pair)

If you are deploying Skype for Business Server on a single standalone or redundant pair of BIG-IP systems, as shown in the alternate logical configuration example on page 4, you must complete these sections of the iApp template:

- Microsoft Skype for Business Edge Virtual Servers - External Interface (if deploying Edge services)
- Edge Server Pools - External Interface (if deploying Edge services)
- Microsoft Skype for Business Reverse Proxy - Reverse Proxy > Forward Reverse Proxy traffic to Skype for Business server(s)
- Microsoft Skype for Business Front End Virtual Servers
- Front End Server Pools
- Microsoft Skype for Business Director Virtual Servers (if deploying Director servers)
- Director server Pools (if deploying Director servers)
- Microsoft Skype for Business Edge Virtual Servers - Internal Interface (if deploying Edge services)
- Edge Server Pools - Internal Interface (if deploying Edge services)

Downloading and importing the Skype for Business iApp

The first task is to download the latest iApp for Microsoft Skype for Business and import it onto the BIG-IP system. You can use this iApp for Skype for Business Server.

To download and import the iApp

1. Open a web browser and go to *https://devcentral.f5.com/codeshare/microsoft-skype-for-business-server-2015*
2. Download the **f5.microsoft_skype_server_2015.<latest-version>.zip** file to a location accessible from your BIG-IP system.
3. Extract (unzip) the **f5.microsoft_skype_server_2015.<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.

Getting started with the Skype for Business iApp

To begin the 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 **Skype-2015_.**
5. From the **Template** list, select **f5.microsoft_skype_server_2015.<latest version>**. The iApp template opens.

Advanced options

If you select **Advanced** from the **Template Selection** list at the very top of the template, you see Device and Traffic Group 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. To use the Device and Traffic Group features, you must have already configured Device and Traffic Groups before running the iApp. For more information on Device Management, see the BIG-IP system documentation.

1. ***Device Group***
   To select a specific Device Group, clear the **Device Group** check box and then select the appropriate Device Group from the list.

2. ***Traffic Group***
   To select a specific Traffic Group, clear the **Traffic Group** check box and then select the appropriate Traffic Group from the list.

Inline help

At the bottom of the Welcome section, the iApp template asks about inline help text.

1. ***Do you want to see inline help?***
   Select whether you want to see informational and help messages inline throughout the template, or if you would rather hide this inline help. If you are unsure, we recommend having the iApp display all inline help.
   Important and critical notes are always shown, no matter which selection you make.

   - **Yes, show inline help text**
     Select this option to see all available inline help text.
   - **No, do not show inline help**
     If you are familiar with this iApp template, or with the BIG-IP system in general, select this option to hide the inline help text.

Configuring the iApp for Skype for Business Front End Servers

This group of questions gathers information for the virtual servers for the Skype for Business Front End Services.

Microsoft Skype Server Front End Virtual Server Questions

Use this section for configuring the iApp for Front End servers.

1. *Are you deploying this system for internal Front End services?*
   Select whether you are deploying the BIG-IP system Skype for Business Front End services at this time.

   - **No, do not deploy this BIG-IP system for Front End services**
     Select this option if you are not deploying the BIG-IP system for Front End Servers at this time. You can always re-enter the template at a later time to add Front End Servers to the deployment.
   - **Yes, deploy this BIG-IP system for Front End services**
     Select this option if you are deploying the BIG-IP system for Front End services.

     a. *Are you deploying this system for all Skype Front End services?*
     Select whether you want to deploy this BIG-IP system for all Skype Front End services or if you are using DNS Load Balancing for non-HTTP Skype Front End services.

     - **Yes, deploy this BIG-IP system for all services**
       Select this option if you are deploying the BIG-IP system for all services, including non-HTTP services such as Skype Server Mediation services.
     - **No, I am using DNS Load Balancing for non-HTTP services**
       Select this option if you are using DNS load balancing for the Front End services that do not use HTTP. If you specify you are using DNS load balancing for non-HTTP services, the system only creates virtual servers for the HTTP-based Front End services.

     b. *What IP address do you want to use for the Front End virtual server?*
     This is the address clients use to access Skype for Business (or a FQDN will resolve to this address). The BIG-IP system will create multiple virtual servers using this address on different ports for the different Front End Services.

     c. *How have you configured routing on your Skype Front End servers?*
     Select whether the Front End servers have a route through the BIG-IP system or not. If the Skype Front End Servers do not have a route back for clients through the BIG-IP system, (i.e. if they do not use the BIG-IP system as the default gateway), the BIG-IP system uses Secure Network Address Translation (SNAT) to translate the client's source address to an address configured on the BIG-IP system.

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

     - **Servers do not have a route to clients through the BIG-IP system**
       Select this option if your Skype Front End Servers do not have a route back to Skype clients through this BIG-IP system.

       a. *How many connections do you expect to each Front End server?*
       Select whether you expect more than 64,000 concurrent connections to each Skype Front End server.

       - **Fewer than 64,000 concurrent connections per server**
         Select this option if you expect fewer than 64,000 concurrent connections per Skype Front End server. With this option, the system applies SNAT Auto Map, which does not require any additional IP addresses, as the system uses an existing self IP address for translation.
       - **More than 64,000 concurrent connections per server**
         Select this option if you expect more than 64,000 connections at one time to each server. With this option, the iApp creates a SNAT Pool, for which you need one IP address for each 64,000 concurrent connections you expect.

         a. *What are the IP addresses you want to use for the SNAT pool?*
         Specify one otherwise unused IP address for every 64,000 concurrent connections, or fraction thereof. Click **Add** for additional rows.

         Important *If you choose more than 64,000 connections, but do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per server is reached, new requests fail.*

     - **Servers have a route to clients through the BIG-IP system**
       Select this option if you have configured a route on the system for traffic from the Front End servers to Skype clients.

d. *On which VLAN(s) should internal Front End traffic be enabled?*
Specify the VLANs from which the BIG-IP system should accept internal Front End traffic. This optional feature can provide an additional layer of security, as you can allow traffic only from the VLANs you choose. The VLAN objects must already be configured on this BIG-IP system before you can select them.

By default, all VLANs configured on the system appear in the Selected (allowed) box. If you do not move any VLANs, the BIG-IP system accepts traffic from all VLANs. Use the Move buttons (<<) and (>>) to adjust list membership.

e. *Have you enabled a hardware load balancing monitoring port on your Front End Servers?*
Specify whether you have enabled a hardware load balancing monitoring port (the default is 5060) on your Front End Servers.

- **No, a hardware load balancing port is not enabled**
Select this option if you have not enabled a hardware load balancing port. No further information is needed.
- **Yes, a hardware load balancing port is enabled**
Select this option if you have enabled a hardware load balancing port. You must answer the following question about which port you are using.

a. *What port have you enabled?*
Specify the monitoring port you are using for hardware load balancing. The default is **5060**.

f. *Are you using Microsoft Skype Server Mediation services?*
This question only appears if you specified you are deploying the BIG-IP system for all Skype Front End services.

Choose whether you are deploying Mediation Services at this time. Skype Mediation services are a necessary component for implementing Enterprise Voice and dial-in conferencing.

- **No, this deployment does not use Mediation servers**
Select this option if you are not deploying the BIG-IP system for Mediation services at this time. You can always re-enter the template at a later time to add this feature.
- **Yes, this deployment uses Mediation services.**
Select this option if you want the BIG-IP system to support Mediation services.

a. *Are your Mediation Servers separate from the Front End Servers?*
The system needs to know if your Mediation Servers are on different servers than your Front End Servers.

- **No, both services are on the same server(s)**
Select this option if your Mediation Servers are on the same servers as your Front End Servers. No further information is required. Continue with the next section.
- **Yes, each service is on a separate server**
Select this option if you want to deploy the BIG-IP system for separate Mediation Servers. A new section appears after the Front End Server Pools section asking for information about your Mediation Servers.

Front End Server Pools

This group of questions gathers information about your Front End Servers to create the BIG-IP load balancing pool.

1. ***Which load balancing method do you want to use?***
Specify the load balancing method you want the BIG-IP system to use for the Front End Servers. While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (node)**.

2. ***Which Front End servers should be in this pool?***
Type the IP address for each Skype Front End Server. You can optionally add a Connection Limit. Click **Add** to include additional servers. You must add at least one Front End Server here.

Front End Mediation Server Pools

This section appears if you specified you are deploying Mediation Servers and they are on different servers than the Front End Servers.

This group of questions gathers information about your Front End Servers to create the BIG-IP load balancing pool.

1. ***Which load balancing method do you want to use?***
Specify the load balancing method you want the BIG-IP system to use for the Mediation Servers. While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (node)**.

2. ***Which Mediation servers should be in this pool?***
Type the IP address for each Mediation Server. You can optionally add a Connection Limit. Click **Add** to include additional servers. You must add at least one Mediation Server here.

Configuring the iApp for Skype for Business Director Servers

This section of the template asks questions about your Skype Server Director servers

Microsoft Skype Server Director Virtual Server Questions

This group of questions gathers information for the virtual servers for the Skype for Business Director servers. Use this section to deploy Director servers to refer internal clients to their home pools. If deploying Edge services, Director servers also proxy external connections for meeting and phone conferencing simple URLs.

1. ***Are you deploying this system for internal Director services?***
   Specify whether you are deploying the BIG-IP system for Skype Server Director servers at this time.
   - **No, do not deploy this BIG-IP system for Director services**
     Select this option if you are not deploying the BIG-IP system for Director servers at this time. You can always re-enter the template at a later time to add Director servers to the deployment.
   - **Yes, deploy this BIG-IP system for Director services**
     Select this option if you are deploying the BIG-IP system for Director servers.
     a. *What IP address do you want to use for this server?*
        Type the IP address the BIG-IP system will use for the Director server virtual server.
        - **Servers have a route to clients through the BIG-IP system**
          Select this option if you have configured a route on the BIG-IP system for traffic coming from the Director servers back to Skype clients.

Director Server Pools

This section only appears if you specified you are deploying Director servers.

This group of questions gathers information about your Director servers to create the BIG-IP load balancing pool.

1. ***Which load balancing method do you want to use?***
   Specify the load balancing method you want the BIG-IP system to use for the Director servers. While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (node)**.
2. ***Which Director servers should be in this pool?***
   Type the IP address for each Director server. You can optionally add a Connection Limit. Click **Add** to include additional servers. You must add at least one Director server here.

Configuring the iApp for Skype for Business Edge Servers - External Interface

This section of the template asks questions about your Skype for Business Server Edge Servers - External Interface. This includes the Access, A/V, and Web Conferencing services.

Important *Be sure to see* *Modifying the iApp configuration on page 27* *for an important change to the virtual server.*

You must provision one unique, publicly routable IP address for each BIG-IP virtual server you create here, plus an additional publicly routable IP address per Edge Server for each Edge service you are deploying. For example, if you are deploying all three services on two Edge Servers, you need to provision nine unique, publicly routable IP addresses.

Microsoft Skype Server Edge Virtual Servers - External Interface

This group of questions gathers information for the virtual servers for the Edge Servers - External Interface.

1. *Are you deploying this system for external Edge services?*
   The first question in this section asks if you are deploying Edge Servers - External Interface at this time. Select **Yes** from the list if you are deploying Edge Servers - External Interface. The Edge Server External Interface options appear.
   - **No, do not deploy this BIG-IP system for external Edge services**
     Select this option if you are not deploying the BIG-IP system for the Edge Server - External Interface at this time. You can always re-enter the template at a later time to add this option to the deployment.
   - **Yes, deploy this BIG-IP system for external Edge services**
     Select this option if you are deploying the BIG-IP system for the Edge Server - External Interface.
     a. *How have you configured routing on your Skype Edge servers?*
        Select whether the Edge servers have a route through the BIG-IP system or not. If the servers do not have a route back for clients through the BIG-IP system, (i.e. if they do not use the BIG-IP system as the default gateway), the BIG-IP system uses Secure Network Address Translation (SNAT) to translate the client's source address to an address configured on the BIG-IP system.

        If you indicate that the Edge servers do have a route back to the clients through the BIG-IP system, the BIG-IP system does not translate the client's source address; in this case, you must make sure that the system is configured as the gateway to the client networks (usually the default gateway) on the servers.
        - **Servers do not have a route to clients through the BIG-IP system**
          Select this option if your servers do not have a route back to Skype clients through this BIG-IP system.
          a. *How many connections do you expect to each Edge server?*
             Select whether you expect more than 64,000 concurrent connections to each server.
             - **Fewer than 64,000 concurrent connections per server**
               Select this option if you expect fewer than 64,000 concurrent connections per server. With this option, the system applies SNAT Auto Map, which does not require any additional IP addresses, as the system uses an existing self IP address for translation.
             - **More than 64,000 concurrent connections per server**
               Select this option if you expect more than 64,000 connections at one time to each server. With this option, the iApp creates a SNAT Pool, for which you need one IP address for each 64,000 concurrent connections you expect.
               a. *What are the IP addresses you want to use for the SNAT pool?*
                  Specify one otherwise unused IP address for every 64,000 concurrent connections, or fraction thereof. Click **Add** for additional rows.

                  Important *If you choose more than 64,000 connections, but do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per server is reached, new requests fail.*
        - **Servers have a route to clients through the BIG-IP system**
          Select this option if you have configured a route on the BIG-IP system for traffic coming from the Edge servers back to Skype clients.

b. *On which VLAN(s) should external Edge traffic be enabled?*
Specify the VLANs from which the BIG-IP system should accept external Edge traffic. This optional feature can provide an additional layer of security, as you can allow traffic only from the VLANs you choose. The VLAN objects must already be configured on this BIG-IP system before you can select them.

By default, all VLANs configured on the system appear in the Selected (allowed) box. If you do not move any VLANs, the BIG-IP system accepts traffic from all VLANs. Use the Move buttons (<<) and (>>) to adjust list membership.

c. *How have you configured Skype Edge services in Skype Topology Builder?*
Choose how you have configured Edge pool services in the Skype Topology Builder. When defining an Edge pool in the Topology Builder, you can specify a single IP address with unique ports for each Edge service, or a unique IP address and FQDN for each service.

- **Skype Edge services use unique FQDNs and IP addresses**
Select this option if each of your Edge services use a unique FQDN and IP address. The system will create a separate virtual server for each service you select in the following questions. If Skype Edge services use a single IP address and FQDN, see *Skype Edge services use a single FQDN and IP address on page 13.*

  a. *What IP address do you want to use for the Access Service virtual server?*
  Type the IP address the BIG-IP system will use for the Edge Servers - External Interface Access Service virtual server. This must be a unique, publicly routable IP address.

  b. *On which port do Edge servers listen for Access Service traffic?*
  Choose the appropriate port on which your Edge Servers listen for Access service traffic. Y

  - **5061**
  Select this option if your Edge servers listen for Access Service traffic on port 5061.
  - **443**
  Select this option if your Edge servers listen for Access Service traffic on port 443.

    a. *Have you enabled federation on port 5061 in the Skype Server Topology?*
    Choose whether you have enabled federation in the Skype Server Topology (on port 5061).

    - **No, federation on port 5061 is not enabled**
    Select this option if you have not enabled federation on port 5061 in your Skype Server Topology. Continue with the next question.
    - **Yes, federation on port 5061 is enabled.**
    Select this option if you have enabled federation in the Skype Server Topology (on port 5061). The iApp creates an additional virtual server for federation.

  c. *Have you enabled federation with XMPP providers on port 5269 in the Skype Topology?*
  Choose whether you have enabled federation with XMPP providers in the Skype Server Topology (on port 5269).

  - **No, federation with XMPP on port 5269 is not enabled**
  Select this option if you have not enabled federation with XMPP on port 5269 in your Skype Server Topology. Continue with the next question.
  - **Yes, federation with XMPP on port 5269 is enabled**
  Select this option if you have enabled federation with XMPP on port 5269 in your Skype Server Topology. The iApp creates an additional virtual server.

  d. *Should the system monitor the internal SIP virtual servers?*
  Select whether you want to create a second monitor to check the health of the internal SIP virtual server. If this server is marked down, the Access Edge Service virtual server will be marked down. This monitor is useful if configuring Skype for Business with BIG-IP GTM in multiple data centers for site resiliency. This is the same monitor described in *Creating a SIP monitor for the Front End servers on page 39, however in this case, the port is 5061.*

  - **No, do not monitor the internal SIP virtual servers**
  Select this option is you do not want to monitor the health of the internal SIP virtual servers. Continue with the next question.
  - **Yes, monitor the internal SIP virtual servers**
  Select this option if you want the BIG-IP system to monitor the health of the internal SIP virtual servers.

    a. *What is the Front End virtual server IP address on the internal BIG-IP LTM?*
    Type the IP address of your internal BIG-IP LTM Front End virtual server. This is the IP address you specified in question 1a *on page 8*.

e. *Are you deploying this system for Web Conferencing services?*
Select whether you are deploying the Edge Server Web Conferencing Service at this time.

- **No, do not deploy this BIG-IP system for Web Conferencing services**
Select this option if you are not deploying the Web Conferencing service at this time. You can always re-enter the template at a later time to add the Web Conferencing service to the configuration.
- **Yes, deploy this BIG-IP system for Web Conferencing services**
Select this option if you want to deploy the BIG-IP system for the Web Conferencing service.

a. *What IP address do you want to use for the Web Conferencing service virtual server?*
If you select Yes, a new row appears asking for the publicly routable IP address you want to use for the Web Conferencing virtual server. Type the IP address the BIG-IP system will use for the Edge Servers - External Interface Web Conferencing Service virtual server.

f. *Are you deploying this system for A/V Edge Services?*
Select whether you are deploying the Edge Server A/V service at this time.

- **No, do not deploy this BIG-IP system for the A/V service.**
Select this option if you are not deploying the A/V service at this time. You can always re-enter the template at a later time to add the A/V service to the configuration.
- **Yes, deploy this BIG-IP system for the A/V service**
Select this option if you want to deploy the BIG-IP system for the A/V service.

a. *What IP address do you want to use for the A/V service virtual server?*
If you select Yes, a new row appears asking for the publicly routable IP address you want to use for the A/V virtual server. Type the IP address the BIG-IP system will use for the Edge Servers - External Interface A/V virtual server.

b. *Should the system translate the source address of A/V service connections?*
Select whether you want the BIG-IP system to use SNAT for the A/V service. For optimal performance, we do not recommend using SNAT for A/V traffic.

> ⚠ ***Warning*** *For best performance, F5 recommends against translating the source address (using SNAT) for Skype A/V traffic, as it is optimal if the Edge servers see the IP address of clients for peer-to-peer client communication. If you do select to translate the source address for A/V connections, the system proxies all A/V traffic through the Edge servers.*

- **No, do not translate the source address of A/V connections**
Select this option if you do not want to use SNAT on the BIG-IP system for A/V traffic. We recommend this option for the best performance.
- **Yes, translate the source address of A/V connections**
Select this option if you want the BIG-IP system to use SNAT for A/V traffic. If you selected **No** in question "*a*" in this section, the system uses the same SNAT setting (Auto Map or a SNAT pool) for the A/V service. Otherwise, if you select to SNAT A/V connections in this question, the system uses SNAT Auto Map.
Continue with *Edge Server Pools - External Interface on page 15*

- **Skype Edge services use a single FQDN and IP address**
Select this option if all of your Skype Edge services use a single FQDN and IP address.

a. *What IP address do you want to use for the Skype Edge services virtual servers?*
Type the IP address the BIG-IP system will use for Edge Servers - External Interface virtual servers. This must be a unique, publicly routable IP address. The system will use this address for all of the Edge services you choose in this section.

b. *On which port do Edge servers listen for Access Service traffic?*
If your Edge Servers listen for Access service traffic on a port other than 5061, type the port number in the box.

c. *Have you enabled federation with XMPP providers on port 5269 in the Skype Topology?*
Choose whether you have enabled federation with XMPP providers in the Skype Server Topology (on port 5269).

- **No, federation with XMPP on port 5269 is not enabled**
Select this option if you have not enabled federation with XMPP on port 5269 in your Skype Server Topology. Continue with the next question.
- **Yes, federation with XMPP on port 5269 is enabled**
Select this option if you have enabled federation with XMPP on port 5269 in your Skype Server Topology. The iApp creates an additional virtual server.

d. *Should the system monitor the internal SIP virtual servers?*
Select whether you want to create a second monitor to check the health of the internal SIP virtual server. If this server is marked down, the Access Edge Service virtual server will be marked down. This monitor is useful if configuring Skype with BIG-IP GTM in multiple data centers for site resiliency. This is the same monitor described in *Creating a SIP monitor for the Front End servers on page 39, however in this case, the port is 5061.*

- **No, do not monitor the internal SIP virtual servers**
Select this option is you do not want to monitor the health of the internal SIP virtual servers. Continue with the next question.
- **Yes, monitor the internal SIP virtual servers**
Select this option if you want the BIG-IP system to monitor the health of the internal SIP virtual servers.

  a. *What is the Front End virtual server IP address on the internal BIG-IP LTM?*
  Type the IP address of your internal BIG-IP LTM Front End virtual server. This is the IP address you specified in question 1a *on page 8*.

e. *Are you deploying this system for Web Conferencing services?*
Select whether you are deploying the Edge Server Web Conferencing Service at this time.

- **No, do not deploy this BIG-IP system for Web Conferencing services**
Select this option if you are not deploying the Web Conferencing service at this time. You can always re-enter the template at a later time to add the Web Conferencing service to the configuration.
- **Yes, deploy this BIG-IP system for Web Conferencing services**
Select this option if you want to deploy the BIG-IP system for the Web Conferencing service.

  a. *On which port do Edge servers listen for Web Conferencing service traffic?*
  Specify the port the Web Conferencing service uses in your Edge implementation. The default is 444. The BIG-IP system creates a virtual server on this port with the IP address you specified at the beginning of this section.

f. *Are you deploying this system for A/V Edge Services?*
Select whether you are deploying the Edge Server A/V service at this time.

- **No, do not deploy this BIG-IP system for the A/V service.**
Select this option if you are not deploying the A/V service at this time. You can always re-enter the template at a later time to add the A/V service to the configuration.
- **Yes, deploy this BIG-IP system for the A/V service**
Select this option if you want to deploy the BIG-IP system for the A/V service.

  a. *On which port do Edge servers listen for A/V service traffic?*
  Specify the port the A/V service uses in your Edge implementation; the default is 443. The system creates a virtual server on this port with the IP address you specified at the beginning of this section.

  b. *Should the system translate the source address of A/V service connections?*
  Select whether you want the BIG-IP system to use SNAT for the A/V service. For optimal performance, we do not recommend using SNAT for A/V traffic.

Warning

For best performance, F5 recommends against translating the source address (using SNAT) for Skype A/V traffic, as it is optimal if the Skype Edge servers see the IP address of clients for peer-to-peer client communication. If you do select to translate the source address for A/V connections, the system proxies all A/V traffic through the Skype Edge servers.

- **No, do not translate the source address of A/V connections**
Select this option if you do not want to use SNAT on the BIG-IP system for A/V traffic. We recommend this option for the best performance.
- **Yes, translate the source address of A/V connections**
Select this option if you want the BIG-IP system to use SNAT for A/V traffic. If you selected **No** in question "*a*" in this section, the system uses the same SNAT setting (Auto Map or a SNAT pool) for the A/V service. Otherwise, if you select to SNAT A/V connections in this question, the system uses SNAT Auto Map.

Edge Server Pools - External Interface

This section only appears if you specified you are deploying Edge Servers - External Interface.

This group of questions gathers information about the load balancing pools for the Edge Servers - External Interface services you are deploying. The number of questions in this section is based on your answers in the previous section.

1. ***Which load balancing method do you want to use for the Access edge service?***
   Specify the load balancing method you want the BIG-IP system to use for the Access Edge service. While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (node)**.

2. ***Which Access Edge servers should be in this pool?***
   Type the IP address for each Access Edge Server. Note these addresses should be publicly routable. You can optionally add a Connection Limit. Click **Add** to include additional servers. You must add at least one server here.

3. ***Which load balancing method do you want to use for the Web Conferencing service?***
   Specify the load balancing method you want the BIG-IP system to use for the Web Conferencing service. While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (node)**.

4. ***Which Web Conferencing servers should be in this pool?***
   Type the IP address for each Web Conferencing Edge Server. Note these addresses should be publicly routable. You can optionally add a Connection Limit. Click **Add** to include additional servers. You must add at least one server here.

5. ***Which load balancing method do you want to use for the A/V Edge service?***
   Specify the load balancing method you want the BIG-IP system to use for the A/V Edge service. While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (node)**.

6. ***Which A/V servers should be in this pool?***
   Type the IP address for each A/V Edge Server. Note these addresses should be publicly routable. You can optionally add a Connection Limit. Click **Add** to include additional servers. You must add at least one server here.

Configuring the iApp for Skype for Business Edge Servers - Internal Interface

This section of the template asks questions about your Edge Servers - Internal Interface deployment. Use this section to deploy Skype for Business Internal Edge services for internally-sourced client connections to external resources.

Microsoft Skype Server Edge Virtual Servers - Internal Interface

This group of questions gathers information for the virtual servers for the Edge Servers - Internal Interface.

1. ***Are you deploying this system for internal Edge services?***
   The first question in this section asks if you are deploying Edge Servers - Internal Interface at this time.

   - **No, do not deploy this BIG-IP system for internal Edge services**
     Select this option if you are not deploying the BIG-IP system for the Edge Server - Internal Interface at this time. You can always re-enter the template at a later time to add this option to the deployment.
   - **Yes, deploy this BIG-IP system for internal Edge services**
     Select this option if you are deploying the BIG-IP system for the Edge Server - Internal Interface.

     a. *What IP address do you want to use for this virtual server?*
     Type the IP address the BIG-IP system will use for the Edge Servers - Internal Interface virtual server.

     b. *How have you configured routing on your Skype Edge servers?*
     Select whether the Edge servers - Internal Interface have a route through the BIG-IP system to internal application clients.

     Note *Note that for the Edge Internal Interface, the default is **Yes**, as typically the internal interface has a route back to the BIG-IP system.*

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

     If you do select Yes from the list, the following question about 64,000 users does not appear.

     - **Servers have a route to internal clients through the BIG-IP system**
       Select this option if the Edge Servers have a route back to internal application clients via this BIG-IP system. No further information is necessary.

     - **Servers do not have a route to clients through the BIG-IP system**
       Select this option if your servers do not have a route back to internal application clients through this BIG-IP system.

       a. *How many connections do you expect to each Edge server?*
       Select whether you expect more than 64,000 concurrent connections to each server.

       - **Fewer than 64,000 concurrent connections per server**
         Select this option if you expect fewer than 64,000 concurrent connections per server. With this option, the system applies SNAT Auto Map, which does not require any additional IP addresses, as the system uses an existing self IP address for translation.
       - **More than 64,000 concurrent connections per server**
         Select this option if you expect more than 64,000 connections at one time to each server. With this option, the iApp creates a SNAT Pool, for which you need one IP address for each 64,000 concurrent connections you expect.

         a. *What are the IP addresses you want to use for the SNAT pool?*
         Specify one otherwise unused IP address for every 64,000 concurrent connections, or fraction thereof. Click **Add** for additional rows.

         Important *If you choose more than 64,000 connections, but do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per server is reached, new requests fail.*

     c. *On which VLAN(s) should internal Edge traffic be enabled?*
     Specify the VLANs from which the BIG-IP system should accept internal Edge traffic. This optional feature can provide an additional layer of security, as you can allow traffic only from the VLANs you choose. The VLAN objects must already be configured on this BIG-IP system before you can select them.

By default, all VLANs configured on the system appear in the Selected (allowed) box. If you do not move any VLANs, the BIG-IP system accepts traffic from all VLANs. Use the Move buttons (<<) and (>>) to adjust list membership.

Edge Server Pools - Internal Interface

This section only appears if you specified you are deploying Edge Servers - Internal Interface.

This group of questions gathers information about the load balancing pools for the Edge Servers - Internal Interface services you are deploying.

1. ***Which load balancing method do you want to use?***
   Specify the load balancing method you want the BIG-IP system to use for the Edge Servers - Internal Interface pool. While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (node)**.

2. ***Which Edge servers should be in this pool?***
   Type the IP address for each internal Edge server. You can optionally add a Connection Limit. Click **Add** to include additional servers. You must add at least one server here.

Configuring the iApp for Skype for Business Reverse Proxy

This section of the template asks questions about whether you are deploying the BIG-IP system for Skype for Business web services (reverse proxy). The configuration described in this section eliminates the need for a separate reverse proxy server in your Skype for Business environment. If you choose to configure the iApp for reverse proxy traffic, you have three options:

- ***Forward reverse proxy client traffic to another BIG-IP system***
  Select this option 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.
- ***Forward reverse proxy client traffic to Skype server(s)***
  Select this option if you are using the BIG-IP system to serve as a reverse proxy for Skype Web Services as described in the previous scenario, and are using a single BIG-IP device (or redundant pair). Skype Web Services traffic is forwarded directly to Front End or Director servers.
- ***Receive reverse proxy traffic from another BIG-IP system***
  Select this option to have the BIG-IP system create internal virtual servers to receive Skype Web Services traffic from a reverse proxy server or external BIG-IP LTM and forward it to the Front End servers. If deploying Director services, requests for simple URLs are forwarded to the Director servers. If deploying a reverse proxy server, such as Microsoft Forefront TMG, configure the proxy publishing rules to forward traffic to the IP addresses of the BIG-IP virtual servers created here.

Select the appropriate option (including choosing not deploy the system for reverse proxy services) from question #1.

Note *If you are upgrading this template from a previous version of the Lync iApp, this template does not save the inputs from the previous Reverse Proxy section(s).*

1. ***Are you deploying this BIG-IP system for Skype web services (reverse proxy) ?***
   The first question in this section asks if you deployed a reverse proxy as part of your Skype for Business Edge topology.
   - **No, do not deploy this BIG-IP system for reverse proxy services**
     Select this option if you are not deploying a reverse proxy at this time. You can always re-enter the template at a later time to add the reverse proxy configuration to the deployment. Continue with the next section.
   - **Forward reverse proxy traffic to another BIG-IP system**
     Select this option to have the iApp create BIG-IP virtual servers to receive external Skype web services traffic and forward it directly to the Skype Front End/Director servers.

     a. Do you want to forward reverse proxy traffic to Director servers?
     Choose whether you want to forward reverse proxy traffic to the Skype Director servers.
     - **Yes, forward reverse proxy traffic to Director servers**
       Select this option if you want the system to forward reverse proxy traffic to the Skype Director servers. The system creates an addition virtual server for this traffic.
     - **No, do not forward reverse proxy traffic to Director servers**
       Select this option if you do not want the BIG-IP system to forward reverse proxy traffic to the Director servers.

     b. Do the pool members (Skype servers or internal BIG-IP) have a route back to application clients via this BIG-IP system?
     If the Internal BIG-IP system does not have a route back for clients through this BIG-IP system, this BIG-IP system uses Secure Network Address Translation (SNAT) to translate the client's source address to an address configured on the BIG-IP system.

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

     We recommend choosing **No** from the list because it does not require you to configure routing manually.
     - **Pool members do not have a route to clients through the BIG-IP system**
       Select this option if the internal BIG-IP system does not have a route back to the application clients through this BIG-IP system.

       a. How many connections do you expect to the virtual server?
       Select whether you expect more than 64,000 concurrent connections to each server.

- **Fewer than 64,000 concurrent connections to the virtual server**
  Select this option if you expect fewer than 64,000 concurrent connections to the virtual server. With this option, the system applies SNAT Auto Map, which does not require any additional IP addresses, as the system uses an existing self IP address for translation.
- **More than 64,000 concurrent connections to the virtual server**
  Select this option if you expect more than 64,000 connections at one time to the virtual server. With this option, the iApp creates a SNAT Pool, for which you need one IP address for each 64,000 concurrent connections you expect.

  a. *What are the IP addresses you want to use for the SNAT pool?*
  Specify one otherwise unused IP address for every 64,000 concurrent connections, or fraction thereof. Click **Add** for additional rows.

  Important *If you choose more than 64,000 connections, but do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per server is reached, new requests fail.*

- **Pool members have a route to clients through the BIG-IP system**
  Select this option if you have configured a route on the internal BIG-IP system for traffic to pass from the servers back to the application clients through this BIG-IP system.

c. *On which VLAN(s) should reverse proxy traffic be enabled?*
Specify the VLANs from which the BIG-IP system should accept reverse proxy traffic. This optional feature can provide an additional layer of security, as you can allow traffic only from the VLANs you choose. The VLAN objects must already be configured on this BIG-IP system before you can select them.

By default, all VLANs configured on the system appear in the Selected (allowed) box. If you do not move any VLANs, the BIG-IP system accepts traffic from all VLANs. Use the Move buttons (<<) and (>>) to adjust list membership.

d. *What IP address do you want to use for the port 443 reverse proxy virtual server?*
Type the unique, publicly routable IP address you want to use for the port 443 reverse proxy virtual server. This virtual server is on port 443.

e. *What is the FQDN of your Skype Front End Web Services pool?*
Type the FQDN you configured in Skype for the External Web Services pool, such as chat.example.com.

f. *What is the FQDN of your Skype Front End Director pool?*
This question only appears if you chose to forward reverse proxy traffic to Director servers in question a.

Type the FQDN you configured in Skype for the Director pool external web services, such as dir.example.com.
When deploying Director Servers, requests for the simple URLs listed in the following questions are forwarded to the Director reverse proxy pool on the internal LTM. If not deploying Director Servers, all requests are forwarded to the internal Front End reverse proxy pool.

g. *What is the simple URL for meetings?*
Type the Meeting Simple URL you specified in your Skype configuration. For example, meet.example.com or www.example.com/meet. Do not use a trailing forward slash in this field.

h. *What is the simple URL for phone access?*
Type the Phone Access Simple URL you specified in your Skype for Business configuration for phone access. For example, dialin.example.com or www.example.com/dialin. Do not use a trailing forward slash in this field.

i. *Do you want to include Skype Mobility services for external clients?*
Select whether you are deploying Skype Mobility services for external clients at this time.

- **No, do not deploy this BIG-IP system for Skype Mobility services**
  Select this option if you do not want to deploy the BIG-IP system for Skype Mobility services for external clients at this time. You can always re-enter the template at a later time to add this functionality to the configuration.
- **Yes, deploy this BIG-IP system for Skype Mobility services**
  Select this option if you want to deploy the BIG-IP system for Skype Mobility services.

  a. *What is the FQDN for external Skype Mobility access?*
  Type the Skype Mobility external URL, such as: skypediscover.example.com

j. *Do you want to create a new client SSL profile for Front End services, or use an existing one?*
Select whether you want the iApp template to create a new client SSL profile for the Front End servers, or if you have already created one on this BIG-IP system for reverse proxy traffic. If you select an existing profile, it must have the appropriate SSL certificate and Key.

Important *If you selected to forward reverse proxy traffic to the Director servers, and plan to use a different Client SSL profile for the Director server traffic, both the Front End and Director Client SSL profiles must be correctly configured for SNI (see the guidance in manual configuration table on page 34) and your clients must support SNI. Otherwise, we recommend using the same SSL profile for both the Front End and Director servers.*

- *Select an existing Client SSL profile*
  If you created a Client SSL profile for this reverse proxy implementation, select it from the list.
- **Create a new Client SSL profile**
  Select this option for the iApp to create a new Client SSL profile using the SSL certificate and key you imported.
  a. *Which SSL certificate do you want to use?*
  Select the SSL certificate you imported for this implementation.
  b. *Which SSL private key do you want to use?*
  Select the associated SSL private key.
  c. *Which intermediate certificate do you want to use?* Advanced
  If your deployment requires an intermediate or chain certificate, select the appropriate certificate from the list. Immediate certificates are intended to create a chain of trust between the CA that signed the certificate and the CA that is already trusted by the recipient of the certificate. This allows the recipient to verify the validity of the certificates presented, even when the signing CA is unknown.

k. *Which client SSL profile do you want to use for Director servers?*
Choose the client SSL profile you want to use for Director server reverse proxy traffic. This question appears because you selected to forward reverse proxy traffic to Director servers. Unless you have a specific need to use a custom client SSL profile or different certificates, we recommend you use the same client SSL profile (which uses the same certificate and key) as the Front End servers.

- *Select an existing Client SSL profile*
  If you created a Client SSL profile for the Director server reverse proxy traffic, select it from the list.

Important *If the profile you created uses a different certificate than the one you are using for the Front End services, it must be configured for SNI, and your clients must support SNI. Otherwise, we recommend using the same SSL profile for both the Front End and Director servers.*

- **Use the same SSL profile as the Front End Servers (recommended)**
  Select this recommended option to have the Director server reverse proxy virtual server use the same client SSL profile as the one you used for the Front End servers. We recommend this option unless you have configured separate certificates for Front End and Director services.
- **Create a new Client SSL profile**
  Select this option for the iApp to create a new Client SSL profile using the SSL certificate and key you imported for the Director servers.
  a. *Which SSL certificate do you want to use?*
  Select the SSL certificate you imported for this implementation.
  b. *Which SSL private key do you want to use?*
  Select the associated SSL private key.

l. *What is the port 4443 virtual server IP address that forwards traffic to the Front End servers?*
Type the IP address of the internal BIG-IP LTM reverse proxy virtual server for external web services that forwards traffic to the Front End servers.

m. *What is the port 4443 virtual server IP address that forwards traffic to the Director Servers?*
Type the IP address of the internal BIG-IP LTM reverse proxy virtual server for external web services that forwards traffic to the Director servers.

This completes the configuration for this scenario. Continue with *Finished on page 27.*

- **Forward reverse proxy traffic client traffic to Skype server(s)**
  Select this option to have the iApp create BIG-IP virtual servers to receive external Skype web services traffic and forward it directly to the Front End/Director servers.

  a. *Do you want to forward reverse proxy traffic to Director servers?*
  Choose whether you want to forward reverse proxy traffic to the Director servers.

  - **Yes, forward reverse proxy traffic to Director servers**
    Select this option if you want the system to forward reverse proxy traffic to the Director servers. The system creates an addition virtual server for this traffic.
  - **No, do not forward reverse proxy traffic to Director servers**
    Select this option if you do not want the BIG-IP system to forward reverse proxy traffic to the Director servers.

  b. *Do the pool members (Skype servers or internal BIG-IP) have a route back to application clients via this BIG-IP system?*
  If the Skype Servers do not have a route back for clients through this BIG-IP system, this BIG-IP system uses Secure Network Address Translation (SNAT) to translate the source address to an address configured on the BIG-IP system.

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

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

  - **Pool members do not have a route to clients through the BIG-IP system**
    Select this option if the Skype Servers do not have a route back to the application clients through this BIG-IP system.

    a. *How many connections do you expect to the virtual server?*
    Select whether you expect more than 64,000 concurrent connections to each server.

    - **Fewer than 64,000 concurrent connections to the virtual server**
      Select this option if you expect fewer than 64,000 concurrent connections to the virtual server. With this option, the system applies SNAT Auto Map, which does not require any additional IP addresses, as the system uses an existing self IP address for translation.
    - **More than 64,000 concurrent connections to the virtual server**
      Select this option if you expect more than 64,000 connections at one time to the virtual server. With this option, the iApp creates a SNAT Pool, for which you need one IP address for each 64,000 concurrent connections you expect.

      a. *What are the IP addresses you want to use for the SNAT pool?*
      Specify one otherwise unused IP address for every 64,000 concurrent connections, or fraction thereof. Click **Add** for additional rows.

      Important *If you choose more than 64,000 connections, but do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per server is reached, new requests fail.*

  - **Pool members have a route to clients through the BIG-IP system**
    Select this option if you have configured a route on the BIG-IP system for traffic to pass from the servers back to the application clients through this BIG-IP system.

  c. *On which VLAN(s) should reverse proxy traffic be enabled?* New
  Specify the VLANs from which the BIG-IP system should accept reverse proxy traffic. This optional feature can provide an additional layer of security, as you can allow traffic only from the VLANs you choose. The VLAN objects must already be configured on this BIG-IP system before you can select them.

  By default, all VLANs configured on the system appear in the Selected (allowed) box. If you do not move any VLANs, the BIG-IP system accepts traffic from all VLANs. Use the Move buttons (<<) and (>>) to adjust list membership.

  d. *What IP address do you want to use for the port 443 reverse proxy virtual server?*
  Type the unique, publicly routable IP address you want to use for the port 443 reverse proxy virtual server. This virtual server is on port 443.

e. *What is the FQDN of your Skype Front End Web Services pool?*
Type the FQDN you configured in Skype for the External Web Services pool, such as chat.example.com.

f. *What is the FQDN of your Skype Front End Director pool?*
This question only appears if you chose to forward reverse proxy traffic to Director servers in question a.

Type the FQDN you configured in Skype for Business for the Director pool external web services, such as dir.example.com. When deploying Director Servers, requests for the simple URLs listed in the following questions are forwarded to the Director reverse proxy pool on the internal LTM. If not deploying Director Servers, all requests are forwarded to the internal Front End reverse proxy pool.

g. *What is the simple URL for meetings?*
Type the Meeting Simple URL you specified in your Skype for Business configuration. For example, meet.example.com or www.example.com/meet. Do not use a trailing forward slash in this field.

h. *What is the simple URL for phone access?*
Type the Phone Access Simple URL you specified in your Skype for Business configuration for phone access. For example, dialin.example.com or www.example.com/dialin. Do not use a trailing forward slash in this field.

i. *Do you want to include Skype Mobility services for external clients?*
Select whether you are deploying Skype Mobility services for external clients at this time.

- **No, do not deploy this BIG-IP system for Skype Mobility services**
  Select this option if you do not want to deploy the BIG-IP system for Skype Mobility services for external clients at this time. You can always re-enter the template at a later time to add this functionality to the configuration.
- **Yes, deploy this BIG-IP system for Skype Mobility services**
  Select this option if you want to deploy the BIG-IP system for Skype Mobility services.
  a. *What is the FQDN for external Skype Mobility access?*
  Type the Skype Mobility external URL, such as: skypediscover.example.com

j. *Do you want to create a new client SSL profile for Front End services, or use an existing one?*
Select whether you want the iApp template to create a new client SSL profile for the Front End servers, or if you have already created one on this BIG-IP system for reverse proxy traffic. If you select an existing profile, it must have the appropriate SSL certificate and Key.

> **Important** *If you selected to forward reverse proxy traffic to the Director servers, and plan to use a different Client SSL profile for the Director server traffic, both the Front End and Director Client SSL profiles must be correctly configured for SNI (see the guidance in manual configuration table on page 34) and your clients must support SNI. Otherwise, we recommend using the same SSL profile for both the Front End and Director servers.*

- *Select an existing Client SSL profile*
  If you created a Client SSL profile for this reverse proxy implementation, select it from the list.
- **Create a new Client SSL profile**
  Select this option for the iApp to create a new Client SSL profile using the SSL certificate and key you imported.
  a. *Which SSL certificate do you want to use?*
  Select the SSL certificate you imported for this implementation.
  b. *Which SSL private key do you want to use?*
  Select the associated SSL private key.
  c. *Which intermediate certificate do you want to use?* Advanced
  If your deployment requires an intermediate or chain certificate, select the appropriate certificate from the list. Immediate certificates are intended to create a chain of trust between the CA that signed the certificate and the CA that is already trusted by the recipient of the certificate. This allows the recipient to verify the validity of the certificates presented, even when the signing CA is unknown.

k. *Which client SSL profile do you want to use for Director servers?*
Choose the client SSL profile you want to use for Director server reverse proxy traffic. This question appears because you selected to forward reverse proxy traffic to Director servers. Unless you have a specific need to use a custom client SSL profile or different certificates, we recommend you use the same client SSL profile (which uses the same certificate and key) as the Front End servers.

- *Select an existing Client SSL profile*
  If you created a Client SSL profile for the Director server reverse proxy traffic, select it from the list.

  Important *If the profile you created uses a different certificate than the one you are using for the Front End services, it must be configured for SNI, and your clients must support SNI. Otherwise, we recommend using the same SSL profile for both the Front End and Director servers.*

- **Use the same SSL profile as the Front End Servers (recommended)**
  Select this recommended option to have the Director server reverse proxy virtual server use the same client SSL profile as the one you used for the Front End servers. We recommend this option unless you have configured separate certificates for Front End and Director services.
- **Create a new Client SSL profile**
  Select this option for the iApp to create a new Client SSL profile using the SSL certificate and key you imported for the Director servers.
  a. *Which SSL certificate do you want to use?*
     Select the SSL certificate you imported for this implementation.
  b. *Which SSL private key do you want to use?*
     Select the associated SSL private key.

l. *Which Front End servers should receive web services traffic?*
Type the IP address(es) of each Front End server that should receive web services traffic. Click **Add** to include additional Front End servers. You can optionally specify a Connection Limit for each server.

m. *Which Director servers should receive web services traffic*
This question only appears if you chose to forward reverse proxy traffic to Director servers in question a.

Type the IP address(es) of each Director server that should receive web services traffic. Click **Add** to include additional Director servers. You can optionally specify a Connection Limit for each server.

This completes the configuration for this scenario. Continue with *Finished on page 27.*

- **Receive reverse proxy traffic from another BIG-IP system**
  Select this option if you want to configure this system to receive reverse proxy traffic from another BIG-IP system. The system creates BIG-IP virtual servers to receive Skype web services traffic from a reverse proxy server or external BIG-IP and forward it to the Front End/Director servers.

  If deploying a third-party reverse proxy server, such as Microsoft Forefront TMG, configure the proxy publishing rules to forward traffic to the IP addresses of the BIG-IP virtual servers you create here. If using BIG-IP LTM to receive the external connections, specify these IP addresses in the Reverse Proxy External Interface section.

  a. *Do you want to forward reverse proxy traffic to Director servers?*
     Choose whether you want to forward reverse proxy traffic to the Director servers.
     - **Yes, forward reverse proxy traffic to Director servers**
       Select this option if you want the system to forward reverse proxy traffic to the Director servers. The system creates an addition virtual server for this traffic.
     - **No, do not forward reverse proxy traffic to Director servers**
       Select this option if you do not want the BIG-IP system to forward reverse proxy traffic to the Director servers.

  b. *Do the pool members (Skype servers or internal BIG-IP) have a route back to application clients via this BIG-IP system?*
     If the Internal BIG-IP system does not have a route back for clients through this BIG-IP system, this BIG-IP system uses Secure Network Address Translation (SNAT) to translate the client's source address to an address configured on the BIG-IP system.

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

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

- **Pool members do not have a route to clients through the BIG-IP system**
  Select this option if the internal BIG-IP system does not have a route back to the application clients through this BIG-IP system.

  a. *How many connections do you expect to the virtual server?*
  Select whether you expect more than 64,000 concurrent connections to each server.

  - **Fewer than 64,000 concurrent connections to the virtual server**
    Select this option if you expect fewer than 64,000 concurrent connections to the virtual server. With this option, the system applies SNAT Auto Map, which does not require any additional IP addresses, as the system uses an existing self IP address for translation.
  - **More than 64,000 concurrent connections to the virtual server**
    Select this option if you expect more than 64,000 connections at one time to the virtual server. With this option, the iApp creates a SNAT Pool, for which you need one IP address for each 64,000 concurrent connections you expect.

    a. *What are the IP addresses you want to use for the SNAT pool?*
    Specify one otherwise unused IP address for every 64,000 concurrent connections, or fraction thereof. Click **Add** for additional rows.

    Important *If you choose more than 64,000 connections, but do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per server is reached, new requests fail.*

- **Pool members have a route to clients through the BIG-IP system**
  Select this option if you have configured a route on the internal BIG-IP system for traffic to pass from the servers back to the application clients through this BIG-IP system.

c. *On which VLAN(s) should reverse proxy traffic be enabled?* New
Specify the VLANs from which the BIG-IP system should accept reverse proxy traffic. This optional feature can provide an additional layer of security, as you can allow traffic only from the VLANs you choose. The VLAN objects must already be configured on this BIG-IP system before you can select them.

By default, all VLANs configured on the system appear in the Selected (allowed) box. If you do not move any VLANs, the BIG-IP system accepts traffic from all VLANs. Use the Move buttons (<<) and (>>) to adjust list membership.

d. *What IP address do you want to use for the Front End port 4443 reverse proxy virtual server?*
Type the unique IP address you want to use for the port 4443 reverse proxy virtual server. This virtual server is on port 4443.

e. *What IP address do you want to use for the Director port 4443 reverse proxy virtual server?*
This question only appears if you chose to forward reverse proxy traffic to Director servers in question a.

Type the unique IP address you want to use for the port 4443 reverse proxy virtual server. This virtual server is on port 4443.

f. *Do you want to create a new client SSL profile for Front End services, or use an existing one?*
Select whether you want the iApp template to create a new client SSL profile for the Front End servers, or if you have already created one on this BIG-IP system for reverse proxy traffic. If you select an existing profile, it must have the appropriate SSL certificate and Key.

- *Select an existing Client SSL profile*
  If you created a Client SSL profile for this reverse proxy implementation, select it from the list.
- **Create a new Client SSL profile**
  Select this option for the iApp to create a new Client SSL profile using the SSL certificate and key you imported.

  a. *Which SSL certificate do you want to use?*
  Select the SSL certificate you imported for this implementation.

  b. *Which SSL private key do you want to use?*
  Select the associated SSL private key.

  c. *Which intermediate certificate do you want to use?* Advanced
  If your deployment requires an intermediate or chain certificate, select the appropriate certificate from the list. Immediate certificates are intended to create a chain of trust between the CA that signed the certificate and the CA that is already trusted by the recipient of the certificate. This allows the recipient to verify the validity of the certificates presented, even when the signing CA is unknown.

g. *Which client SSL profile do you want to use for Director servers?*
Choose the client SSL profile you want to use for Director server reverse proxy traffic. This question appears because you selected to forward reverse proxy traffic to Director servers. Unless you have a specific need to use a custom client SSL profile or different certificates, we recommend you use the same client SSL profile (which uses the same certificate and key) as the Front End servers.

- *Select an existing Client SSL profile*
  If you created a Client SSL profile for the Director server reverse proxy traffic, select it from the list. Unless you have specific requirements, we recommend using the same certificate and key used for the Front End services.
- **Use the same SSL profile as the Front End Servers (recommended)**
  Select this recommended option to have the Director server reverse proxy virtual server use the same client SSL profile as the one you used for the Front End servers. We recommend this option unless you have configured separate certificates for Front End and Director services.
- **Create a new Client SSL profile**
  Select this option for the iApp to create a new Client SSL profile using the SSL certificate and key you imported for the Director servers.

> **Important**
>
> *If you are using a different certificate than the one you are using for the Front End services, it must be configured for SNI, and your clients must support SNI. Otherwise, we recommend using the same SSL profile for both the Front End and Director servers.*

a. *Which SSL certificate do you want to use?*
Select the SSL certificate you imported for this implementation.

b. *Which SSL private key do you want to use?*
Select the associated SSL private key.

h. *Which Front End servers should receive web services traffic*
Type the IP address(es) of each Front End server that should receive web services traffic. Click **Add** to include additional Front End servers. You can optionally specify a Connection Limit for each server.

i. *Which Director servers should receive web services traffic*
This question only appears if you chose to forward reverse proxy traffic to Director servers in question a.

Type the IP address(es) of each Director server that should receive web services traffic. Click **Add** to include additional Director servers. You can optionally specify a Connection Limit for each server.

Advanced Firewall Manager (BIG-IP AFM)

If you chose to deploy the BIG-IP system for External Edge services, or are forwarding reverse proxy client traffic another BIG-IP system or the Skype servers, you have the option of using the BIG-IP Advanced Firewall Manager to protect the deployment. For more information on configuring BIG-IP AFM, see *http://support.f5.com/kb/en-us/products/big-ip-afm.html*, and then select your version.

1. ***Do you want to use BIG-IP AFM to protect your application?***
   Choose whether you want to use BIG-IP AFM, F5's network firewall, to secure this Skype for Business deployment. If you choose to use BIG-IP AFM, you can restrict access to the Skype virtual server to a specific network or IP address. See the BIG-IP AFM documentation for specific details on configuring AFM.

   - **No, do not use Application Firewall Manager**
     Select this option if you do not want to enable BIG-IP AFM at this time. You can always re-enter the template at a later date to enable BIG-IP AFM. Continue with the next section.
   - *Select an existing AFM policy from the list*
     If you already created a BIG-IP AFM policy for this implementation, select it from the list. Continue with ***c***.
   - **Yes, use F5's recommended AFM configuration**
     Select this option if you want to enable BIG-IP AFM using F5's recommended configuration.

a. *Do you want to restrict access to your application by network or IP address?*
Choose whether you want to restrict access to the Skype implementation via the BIG-IP virtual server.

- **No, do not restrict source addresses (allow all sources)**
By default, the iApp configures the AFM to accept traffic destined for the Skype virtual server from all sources. If you do not have a need to restrict access to the virtual server, leave this option selected and then continue with ***b***.
- **Restrict source addresses**
Select this option if you want to restrict access to the Skype virtual server by IP address or network address. You specify addresses that should be allowed to access the application, and can also choose to deny access to external resources in the subsequent question.

    a. *What IP or network addresses should be allowed to access your application?*
    Specify the IP address or network access that should be allowed access to the Skype virtual server. You can specify a single IP address, a list of IP addresses separated by spaces (not commas or other punctuation), a range of IP addresses separated by a dash (for example **`192.0.2.10-192.0.2.100`**), or a single network address, such as **`192.0.2.200/24`**.

b. *How do you want to control access to your application from sources with a low reputation score?*
The BIG-IP AFM uses an IP intelligence database to categorize IP addresses coming into the system. Choose what you want the system to do for sources that are attempting to access the Skype virtual server with a low reputation score. For more information, see the BIG-IP AFM documentation.

> ***Important*** *You must have an active IP Intelligence license for this feature to function. See https://f5.com/products/modules/ip-intelligence-services for information.*

- **Allow all sources regardless of reputation**
Select this option to allow all sources, without taking into consideration the reputation score.
- **Reject access from sources with a low reputation**
Select this option to reject access to the Skype virtual server from any source with a low reputation score.
- **Allow but log access from sources with a low reputation**
Select this option to allow access to the Skype virtual server from sources with a low reputation score, but add an entry for it in the logs.

c. *Would you like to stage a policy for testing purposes?*
Choose whether you want to stage a firewall policy for testing purposes. A staged policy allows you to evaluate the effect a policy has on traffic by analyzing the system logs, without actually modifying traffic based on the firewall rules. You must already have a policy on the system in order to select it.

- **Do not apply a staging policy**
Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.
- *Select an existing policy from the list*
If you have already created a firewall policy for this implementation, select it from the list. Only policies that already exist on the system appear in the list. To create a new policy, on the Main tab, click **Security > Network Firewall > Policies**. Specific instructions for creating a firewall policy is outside the scope of this iApp and deployment guide.

d. *Which logging profile would you like to use?*
Choose whether you or not you want to use a logging profile for this AFM implementation. You can configure the BIG-IP system to log detailed information about BIG-IP system Network Firewall events and store those logs on the BIG-IP system or a remote logging server (supports syslog and Splunk). If you want to use a logging profile, we recommend creating one outside this template. The list only contains profiles with Network Firewall enabled.

- **Do not apply a logging profile**
Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.
- *Select an existing logging profile from the list*
If you have already created a logging profile for this implementation, select it from the list. You must create a profile before it is available in the list. To create a logging profile, on the Main tab, click **Security > Event Logs > Logging Profiles**. Specific instructions for creating a logging profile is outside the scope of this iApp and deployment guide. See the online help or the *About Local Logging with the Network Firewall* chapter of the ***BIG-IP Network Firewall: Policies and Implementations*** guide for more information.

Finished

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

Modifying the iApp configuration

If you configured the iApp for Microsoft Skype Server Edge Servers: External Interface, and specified you were deploying the system for A/V Edge services, you must make a change to the virtual server configuration after completing the iApp template.

First, if you have not yet disabled Strict Updates, click **iApps > Application Services** and then click the name of your Skype for Business application service. On the menu, click Properties, and then from the Application Service list, select **Advanced**. In the **Strict Updates** row, clear the box to disable Strict Updates.

Next, click **Local Traffic > Virtual Servers** and then from the list, click the name of the external UDP virtual server on port 3478. From the **Source Port** list, select **Change**, and then click **Update**.

Troubleshooting

Use this section for common issues and troubleshooting steps.

- **Skype for Business clients cannot connect or receive authentication prompts when accessing Microsoft Exchange Autodiscover and EWS through F5 APM**

  When you have deployed BIG-IP APM in front of Microsoft Exchange , Microsoft Skype for Business clients may be unable to successfully query the Autodiscover service or download free/busy information from EWS. Because this is an issue with the BIG-IP system and Exchange, to work around this issue, you must create an iRule to disable APM for these requests in your BIG-IP configuration for Exchange server. For specific instructions, see the Exchange deployment guide, available at *https://f5.com/solutions/deployment-guides/microsoft-exchange-server-2010-and-2013-big-ip-v11*

Creating a forwarding virtual server for Skype Edge server to Skype client communication

When you use F5's recommended configuration for Skype Edge services (includes both manual and iApp template configuration), you must create a forwarding BIG-IP virtual server to accept outbound traffic from the Edge server. Because the Edge server(s) use the BIG-IP self IP address as a default gateway, this BIG-IP virtual server must be configured to allow asymmetric traffic to pass through the BIG-IP LTM when the Edge server is responding to direct Skype client requests.

For this configuration, you must create a Fast L4 Profile and a virtual server. Use the following table for guidance. For information on configuring specific objects, see the online help or BIG-IP documentation.

BIG-IP LTM Object	Non-default settings/Notes	
Fast L4 Profile (*Local Traffic-->Profiles-->Protocol-->Fast L4*)	***Name***	Type a unique name
	Parent Profile	**http**
	Loose Initiation	**Enabled**
	Loose Close	**Enabled**
Virtual Server (*Main tab-->Local Traffic-->Virtual Servers*)	***Name***	Type a unique name
	Type	**Performance Layer 4**
	Destination Address	**0.0.0.0/0**
	Protocol	**All Protocols**
	Protocol Profile (Client)	Select the Fast L4 profile you created
	VLANs and Tunnels	Select appropriate VLAN(s)
	Source Address Translation	**None**
	Address Translation	Clear the **Enabled** box to **disable** Address Translation
	Port Translation	Clear the **Enabled** box to **disable** Port Translation

You must also have configured a network route on the BIG-IP system for forwarding traffic from Skype Edge servers to the client network(s). To configure a BIG-IP route, see **Network > Routes**. For specific information or help, see the BIG-IP documentation or online help.

Next steps

After completing the iApp Template, the BIG-IP Application Services page opens for the application service you just created. To see the list of all the configuration objects created to support Skype for Business Server , on the Menu bar, click **Components**. The complete list of all Skype for Business server 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 Skype for Business Server 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). 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 Skype for Business Server 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 Skype for Business Server configuration objects.

To view object-level statics

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.

Appendix: Manual Configuration table for BIG-IP objects

Because of the complexity of this configuration, we strongly recommend using the iApp template to configure the BIG-IP system for Skype Server . Advanced users extremely familiar with the BIG-IP can use following tables to configure the BIG-IP manually. This first table shows the non-default settings on BIG-IP objects for the Front End Services. BIG-IP pool members (column 2) are each of the 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: Skype for Business 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**	**Skype-http-fe**: Base HTTP parent **Skype-tcp-5061-fe:**[6] Base TCP parent: - **Alias Service Port**: **5061**	***Skype-tcp-fe***: Base TCP Parent profile with **Idle Timeout** set to **1800**	**Skype-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**	**Skype-tcp-monitor-fe**: Base TCP parent with no required changes	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	RPC
Service Port: **443**	Service Port: **443**[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-fe* **and** *Skype-tcp-5061-fe*	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	HTTPS
Service Port: **444**	Service Port: **444**[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-fe*	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	
Service Port: **448**	Service Port: **448**[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-fe*	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	
Service Port: **5061**	Service Port: **5061**[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-fe* *Optional monitor* [3]*:* **Skype-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 *Skype-tcp-monitor-fe*	Use *Skype-tcp-fe*	Use *Skype-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 *Skype-tcp-monitor-fe*	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	Same note as above
Service Port: **5070**[5]	Service Port: **5070**[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-fe*	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	Same note as above
Service Port: **5071**	Service Port: **5071**[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-fe*	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	
Service Port: **5072**	Service Port: **5072**[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-fe*	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	
Service Port: **5073**	Service Port: **5073**[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-fe*	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	

[1] *Use the Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 39*)
[3] For the SIP monitor, additional steps need to be taken on the Microsoft Front-End Servers. See *Creating a SIP monitor for the Front End servers on page 39*
[4] SSL persistence is optional but recommended
[5] These virtual servers are only necessary if deploying Mediation Servers.
[6] This TCP monitor is used to support bringing down pool members when Skype 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 *Skype-tcp-monitor-fe*	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	
Service Port: **5076**	Service Port: **5076**[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-fe*	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	
Service Port: **5080**	Service Port: **5080**[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-fe*	Use *Skype-tcp-fe*	Use *Skype-source-fe*	Yes [2]	
Service Port: **8080**	Service Port: **8080**[1] **Action on Service down: Reject**	Use *Skype-http-fe* **and** *Skype-tcp-5061-fe*[6]	Use *Skype-tcp-fe* **HTTP**: *Skype-fe-http* Base HTTP parent with no optimizations	**Skype-cookie-fe**: Default profile with Type set to **Cookie** persistence.	Yes [2]	

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 39*)
[3] For the SIP monitor, additional steps need to be taken on the Microsoft Front-End Servers. See *Creating a SIP monitor for the Front End servers on page 39*
[4] SSL persistence is optional but recommended
[5] These virtual servers are only necessary if deploying Mediation Servers.
[6] This TCP monitor is used to support bringing down pool members when Skype servers are put into Maintenance Mode.

Configuration table for BIG-IP objects: Skype for Business Director Services

The following table shows the non-default settings on BIG-IP LTM objects for the Director services. The BIG-IP pool members (column 2) for the following table are each of the 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**	**Skype-tcp-monitor-dir**: Base TCP monitor with no required changes **Skype-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 *Skype-tcp-monitor-dir*	Standard TCP	None	Yes [2]	
5061	Service Port: 5061[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-dir*	Standard TCP	None	Yes [2]	SIP over TLS

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 39*)
[3] This TCP monitor is used to support bringing down pool members when Skype servers are put into Maintenance Mode.

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

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

➲ **Note** *When defining an Edge pool in the Topology Builder, you specify a single IP address with unique ports for each Edge service, or a unique IP address and FQDN for each service. If you configured the Topology Builder for separate FQDNs for Web Conferencing and A/V, each Skype for Business 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. If you specified a single IP address and FQDN, you only need one publicly routable IP address on each server in that case.*

Virtual Server port	Pool	Health monitor	Profiles	Persistence profile	SNAT?	Notes
Important: If you configured the Topology Builder for a single FQDN and IP address when defining an Edge pool, use the same IP address for each of the following virtual servers. If you configured unique FQDNs for the web conferencing and A/V services, use a unique IP address for each service.						
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 Topology, and created the Access virtual server on port 443, you must **also** create the virtual server on port 5061.						
443	Service Port: 443[1] **Action on Service down: Reject**	**Skype-tcp-monitor-ext**: Base TCP monitor with no required changes	**TCP**: *Skype-edge-tcp-ext*: Base *tcp* parent profile with **Idle Timeout** set to **1800** **Nagle's Algorithm**: **Disabled**	Source Address Affinity	Yes [2]	
5061 (default for single IP address)	Service Port: 5061[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-ext*	Use *Skype-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 *Skype-tcp-monitor-ext*	Use *Skype-edge-tcp-ext*	Source Address Affinity	Yes [2]	
Web Conferencing Service						
443 (444 for single IP address)	Service Port: 443[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-ext*	Use *Skype-edge-tcp-ext*	Source Address Affinity	Yes [2]	
A/V Service [3]						
443 (default for single IP address)	Service Port: 443[1] **Action on Service down: Reject**	Use *Skype-tcp-monitor-ext*	Use *Skype-edge-tcp-ext*	Source Address Affinity	Not recommended[4]	The A/V Edge external interfaces must have publicly routable IP addresses
3478 (see trouble-shooting on *page 28)*	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]	On the virtual server, the **Source Port** list must be set to **Change**. Add both monitors to the pool. The iCMP monitor ensures a pool member is properly marked down

[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 39*)
[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 Skype Edge Servers.
The BIG-IP pool members (column 2) for the following table are the internal interface of the Edge Servers.

Virtual Server Port	Pool	Health monitor	Profiles	Persistence Profile	SNAT?	Notes
443	Service Port: 443[1] **Action on Service down: Reject**	**Skype-tcp-monitor-int**: Base TCP monitor with no required changes	***TCP:*** *Skype-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 *Skype-tcp-monitor-int*	Use *Skype-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 *Skype-tcp-monitor-int*	Use *Skype-edge-tcp-int*	*Default*: SSL [3] **Timeout** set to **1800** *Fallback*: Source Address Affinity	Yes [2]	

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 39*)
[3] SSL persistence is optional but recommended

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.

There are three options for configuring the BIG-IP LTM when using a reverse proxy. Follow the guidance applicable to your configuration.

- *Receive Reverse Proxy traffic from another BIG-IP system configuration table on page 35*
- *Forward Reverse Proxy client traffic to another BIG-IP system on page 36*
- *Forward Reverse Proxy traffic to Skype for Business Server(s) on page 37*

Receive Reverse Proxy traffic from another BIG-IP system 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	SNAT enabled?
Front End reverse proxy virtual server				
4443	Front End pool members on port 4443[1] **Action on Service down: *Reject***	**Skype-https-4443-fe**: Base HTTPS monitor **Alias Service Port** set to **4443** Other settings optional **Skype-tcp-5061-in-rp[3]:** Base TCP parent: **Alias Service Port**: **5061**	Use *Skype-tcp-fe* **Client SSL**: *Skype-fe-client-ssl*: Base client SSL profile. *Important:* Must use same certificate used by Skype Server. **Server SSL**: Skype-fe-server-ssl: Base server SSL profile with proper certs **HTTP**: *Skype-fe-http* Base HTTP parent profile with no optimizations	Yes [2]

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 39*)
[3] This TCP monitor is used to support bringing down pool members when Skype 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	SNAT enabled?
Front End reverse proxy virtual server				
4443	Director server pool members on port 4443[1] **Action on Service down: Reject**	**Skype-https-4443-fe**: Base HTTPS monitor **Alias Service Port** set to **4443** Other settings optional **Skype-tcp-5061-in-rp[3]:** Base TCP parent: **Alias Service Port**: **5061**	Use *Skype-tcp-fe* **Client SSL**: *Skype-fe-client-ssl*: Base client SSL profile. *Important:* Must use same certificate used by Skype Server. **Server SSL**: Skype-fe-server-ssl: Base server SSL profile with proper certs. **HTTP**: *Skype-fe-http* Base HTTP parent profile with no optimizations	Yes [2]

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 39*)
[3] This TCP monitor is used to support bringing down pool members when Skype servers are put into Maintenance Mode.

NOTE: When deploying an external reverse proxy for Skype web services, F5 recommends either deploying an LTM virtual server to receive external Skype 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.

Forward Reverse Proxy client traffic to another BIG-IP system

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	SNAT?	Other
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**	**Skype-https-4443-fe**: Base HTTPS monitor **Alias Service Port** set to **4443** Other settings optional **Skype-tcp-5061-ex-rp**[3] Base TCP parent: **Alias Service Port**: **5061**	Use *Skype-tcp-fe* **Server SSL**: Skype-fe-server-ssl: Base server SSL profile with proper certs. **HTTP**: *Skype-fe-http* Base HTTP parent profile with no optimizations **Client SSL**: *Skype-fe-client-ssl*: Base client SSL profile. *Important:* Must use same certificate used by Skype Server. **If using Director servers and a unique certificate**[4]: Set the **Server Name** to the FQDN of your Front End web services pool. You must also create a Director Client SSL profile: *Skype-dir-client-ssl*: Base client SSL profile with **Default SSL profile for SNI** set to **Enabled**.	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 *Creating the iRules on page 38*

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 39*)
[3] This TCP monitor is used to support bringing down pool members when Skype servers are put into Maintenance Mode.
[4] If using a unique certificate for the Director servers, the Client SSL profile must be configured for SNI, and your clients must support SNI

Forward Reverse Proxy traffic to Skype for Business Server(s)

Create the following virtual server which will receive external Skype for Business web services traffic and forward it directly to the Front End/Director servers.. 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?	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*: Create a pool with the Front End servers that should receive web services traffic. *Director reverse proxy pool*: If using Director servers, create an additional pool with the Director servers that should receive web services traffic. Both use Service Port **4443**[1] **Action on Service down: Reject**	**Skype-https-4443-fe**: Base HTTPS monitor **Alias Service Port** set to **4443** Other settings optional **Skype-tcp-5061-ex-rp**[3] Base TCP parent: **Alias Service Port**: **5061**	Use *Skype-tcp-fe* **Server SSL**: Skype-fe-server-ssl: Base server SSL profile with proper certs. **HTTP**: *Skype-fe-http* Base HTTP parent profile with no optimizations **Client SSL**: *Skype-fe-client-ssl*: Base client SSL profile. *Important:* Must use same certificate used by Skype Server. **If using Director servers and a unique certificate**[4]: Set the **Server Name** to the FQDN of your Front End web services pool. You must also create a Director Client SSL profile: *Skype-dir-client-ssl*: Base client SSL profile with **Default SSL profile for SNI** set to **Enabled**.	***If using a single BIG-IP LTM only:*** Cookie: Cookie Name set to **MS-WSMAN** Always Send Cookie**: Enabled** Expiration**: 3650 days** **(this profile is optional)**	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 *Creating the iRules on page 38*

[1] Select **Advanced** from the **Configuration** list, and use the *Least Connections (node)* load balancing method
[2] ***Required*** (see *Creating a SNAT on page 39*)
[3] This TCP monitor is used to support bringing down pool members when Skype servers are put into Maintenance Mode.
[4] If using a unique certificate for the Director servers, the Client SSL profile must be configured for SNI, and your clients must support SNI

Creating the iRules

For the external reverse proxy virtual server, you must create an iRule that sends traffic to the proper Skype 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.

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 Front End Server configuration to enable for 5060.

To enable port 5060, use the 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 Skype Server Deployment Wizard on each Front End server to apply the change.

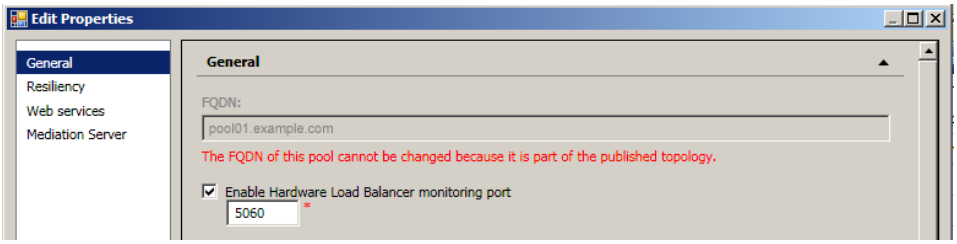

Figure 3: *Editing the General properties*

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 **Skype-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, Skype will respond to SIP requests on that port with a status code of "488" (and "401" if using NTLM authentication) and the 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 a SNAT

A source network address translation (SNAT) allows for inter-server communication and provides the ability to perform certain Skype 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 Skype Server hosts.

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

Use the procedure most applicable for your deployment.

As mentioned in the prerequisites, we typically recommend *Auto Map* for SNAT configuration. With SNAT Auto Map 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 Skype 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 **Skype-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 Skype Servers reside, and then click the Add (**<<**) button.
7. Click the **Finished** button.

Configuring a SNAT for large Skype Server deployments
For large deployments (with 64,000 simultaneous connections), 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 Skype Servers. You must create a SNAT pool for each 64,000 connections (or fraction thereof).

Important *This procedure is only necessary for large deployments. If your Skype deployment has less than 64,000 simultaneous connections, 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 **Skype-snat-pool**.
5. In the **IP Address** box, type in a valid and otherwise-unused address on the subnet containing your Front End servers, and click the **Add** button.
   Repeat this step for each additional address needed. At least one address should be added for each 64,000 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 **Skype-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 **Skype-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 Skype devices reside, and click the Add (<<) button.
14. Click the **Finished** button.

Manually configuring the BIG-IP Advanced Firewall Module to secure your Skype for Business deployment

If you chose to deploy the BIG-IP system for External Edge services, or are forwarding reverse proxy client traffic another BIG-IP system or the Skype servers, you have the option of using the BIG-IP Advanced Firewall Manager to protect the deployment. This section describes how to manually configure BIG-IP AFM to secure your Skype for Business deployment. BIG-IP AFM is particularly useful if you want to only allow access from specific clients or networks. Because this configuration can be complex, we recommend using the iApp template in version 11.6 and later to configure BIG-IP AFM.

Network Firewall settings

When configuring the BIG-IP Advanced Firewall Manager, you may want to configure your BIG-IP system to drop all traffic that you have not specifically allowed with firewall rules. This in known as ***firewall mode***. By default, your BIG-IP system is set to default-accept, or ***ADC mode***. Instructions for configuring your BIG-IP system, and the implications to consider, can be found on AskF5. For example, for BIG-IP v11.5: *http://support.f5.com/kb/en-us/products/big-ip-afm/manuals/product/network-firewall-policies-implementations-11-5-0/1.html*.

If you have licensed IP Intelligence on the BIG-IP system, you can prohibit connections from sources with low reputation scores.

The following instructions cover a basic firewall configuration that is effective for the most common scenario of wanting to allow connections from a single trusted network. If you have complex requirements, such as the need to schedule different policies for different times of the day, or you want to create complicated rule or address lists, consult the BIG-IP AFM documentation. The basic steps for Policy and Rule creation apply to all scenarios.

To configure the BIG-IP AFM to allow connections from a single trusted network

1. Create a Network Firewall Policy:
    a. From the Configuration utility, click **Security > Network Firewall > Policies**, and then click **Create**.
    b. In the **Name** field, type a unique name for the policy, such as **Skype-Policy**.
    c. Click **Finished**.
2. Create a rule to allow authorized hosts or networks to connect:
    a. Click **Security > Network Firewall > Policies**.
    b. Click the name of the policy you just created.
    c. In the Rule section (below the General Properties section), click the **Add** button.
    d. Leave the **Type** list set to Rule.
    e. From the **Order** list, select **First**. The Order list only appears in version 11.5 and later. In 11.4.x, you must reorder the rules from the Policy General Properties page.
    f. In the **Name** field, type a unique name, for instance **Skype-traffic-Allowed**.
    g. Ensure the **State** list is set to **Enabled**.
    h. From the **Protocol** list, select **TCP**. Leave the box to the right of TCP set to **6**.
    i. In the **Source** section, from the **Address/Region** list, select **Specify**.
    You are now able to list the trusted source addresses for your connection.
    In the following example, we will configure a single subnet as trusted.
        - Select **Address**.
        - In the box, type the network address you want to allow, including netmask if more than a single host. Specify a network using CIDR notation, such as **10.0.0.0/24**.
        - Do not configure a source port.
        - Optional: If you want to limit inbound connections to a specific VLAN or Tunnel, from the **VLAN / Tunnel** list, select **Specify**, and then move the VLANs or tunnels that are allowed access to the Selected box.
        - Click **Add**.
        - Repeat these steps for additional hosts or networks. Use **Address List** or **Address Range** when appropriate.
    j. In the **Destination** section, leave the **Address/Region** and **Port** set to **Any**. Because you will be applying your policy to a virtual server that listens only on a single desired address and port, do not specify that information here.

k. If necessary, from the **Action** list, select **Accept**.

l. *Optional:* If you have configured a logging profile and want to log connections, from the **Logging** list, select **Enabled**. Typically, allowed connections do not need to be logged.

m. Click **Finished**.

3. Creating a firewall rule to block all other traffic

The next task is to create a firewall rule to block all other traffic that you have not allowed. Although this is not a required step if your BIG-IP system is set to default deny (**Firewall mode**), it is required in default-accept (**ADC mode**), and is a good practice to always configure such a rule.

a. Click **Security > Network Firewall > Policies**.

b. Click the name of the policy you created in step 1.

c. In the Rule section (below the General Properties section), click the **Add** button.

d. Leave the **Type** list set to **Rule**.

e. Leave the **Order** list, select **Last**.

f. In the **Name** field, type a unique name, for example **Skype-traffic-Prohibited**.

g. Ensure the **State** list is set to **Enabled**.

h. From the **Protocol** list, select **TCP**. Leave the box to the right of TCP set to **6**.

i. In the **Source** section, leave all the lists set to **Any**

j. From the **Action** list, select either **Drop** (to silently discard incoming connections) or **Reject** (to send a Destination Unreachable message to the sender).

k. If you configured a logging profile as described in *Optional: Configuring the BIG-IP system to log network firewall events on page 44*, from the **Logging** list, select **Enabled**.

l. Click **Finished**. You return to the Policy Properties page.

m. On the Policy Properties page, in the Rules section, ensure the rule with the Action of Accept comes before the Drop or Reject rule you just created. If it does not, use the **Reorder** button and drag the rules into the correct order.

4. Apply Your Firewall Policy to your Virtual Server

a. Click **Security > Network Firewall > Active Rules**.

b. In the Rule section (below the General Properties section), click the **Add** button.

c. From the **Context** list, select **Virtual Server**, and then select the virtual server you created for your Skype External Edge or reverse proxy traffic.

d. From the **Type** list, select **Policy**, and then select the firewall policy you created.

e. From the **Policy Type** list, select **Enforced**.

f. Click **Finished**.

Optional: Assigning an IP Intelligence Policy to your Skype virtual server

If you want to restrict access to your Skype virtual server based on the reputation of the remote sender, you can enable and assign an IP Intelligence policy. This requires an IP intelligence license; contact your F5 Sales representative for more information.

It is outside the scope of this document to provide instructions on configuring an IP Intelligence Policy. Full documentation on enabling and configuring the IP Intelligence feature can be found on AskF5. For example, the manual for BIG-IP AFM v11.5 is: *https://support.f5.com/kb/en-us/products/big-ip-afm/manuals/product/network-firewall-policies-implementations-11-5-0/5.html*

After you have enabled and configured an IP Intelligence policy, use the following steps to assign the policy to your virtual server:

To assign the IP intelligence policy to the Skype virtual server

1. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.

2. Click the name of your Skype for Business External Edge or reverse proxy virtual server.
3. From the **Security** menu, choose **Policies**.
4. Next to **IP Intelligence**, select **Enabled**, then select the IP intelligence policy to apply to traffic on the virtual server.
5. Click **Update**. The list screen and the updated item are displayed. The IP Intelligence policy is applied to traffic on the virtual server.

Optional: Configuring the BIG-IP system to log network firewall events

If you are using AFM, you have the option of logging network firewall events to one or more remote syslog servers (recommended) or to log events locally. You can either use an iApp template to create the logging profile, or create the logging profile manually.

For specific information on logging on the BIG-IP system, see the appropriate guide for your version. For example, for 11.5.0:

- Remote High-Speed Logging: *https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos-implementations-11-5-0/22.html*
- Local logging: *https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos-concepts-11-5-0/11.html*

Creating the logging profile using the iApp template

Use this section to create the logging profile using the logging profile iApp template. If you have not already downloaded the iApp template, see *https://devcentral.f5.com/wiki/iApp.F5-Remote-Logging-iApp.ashx*.

To configure the logging profile iApp

1. Log on to the BIG-IP system.
2. On the Main tab, click **iApp > Application Services**.
3. Click **Create**. The Template Selection page opens.
4. In the **Name** box, type a name. In our example, we use **logging-iapp_.**
5. From the **Template** list, select **f5.remote_logging.v*<latest-version>***. The template opens
6. Use the following table for guidance on configuring the iApp template. Questions not mentioned in the table can be configured as applicable for your implementation.

Question	Your selection
Do you want to create a new pool of remote logging servers, or use an existing one?	Unless you have already created a pool on the BIG-IP system for your remote logging servers, select **Create a new pool**.
Which servers should be included in this pool?	Specify the IP addresses of your logging servers. Click **Add** to include more servers.
What port do the pool members use?	Specify the port used by your logging servers, typically **514**.
Do the pool members expect UDP or TCP connections?	**TCP**
Do you want to create a new monitor for this pool, or use an existing one?	Unless you have already created a health monitor for your pool of logging servers, select **Use a simple ICMP (ping) monitor**.
Do your log pool members require a specific log format?	If your logging servers require a specific format, select the appropriate format from the list.

7. Click **Finished**.
8. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.
9. Click the name of your Skype for Business External Edge or reverse proxy virtual server.
10. From the **Security** menu, choose **Policies**.
11. Next to **Log Profile**, select **Enabled**, then select the Logging profile you created.
12. Click **Update**. The list screen and the updated item are displayed

Note *The iApp template creates a log publisher and attaches it to the logging profile. If the publisher does not appear in the BIG-IP Configuration utility (GUI), you can verify the configuration by running the following command from the Traffic Management shell (tmsh):* **`list security log profile <your profile name>`**.

Creating logging profile manually

If you do not want to use the iApp template to create a logging profile, use this section for guidance on configuring the logging profile manually. You must have access to the tmsh command line to use this method.

To manually configure a logging profile

1. Use the following guidance for configuring a health monitor and load balancing pool for the logging servers.

BIG-IP LTM Object	Non-default settings/Notes	
Health Monitor (*Local Traffic -->Monitors*)	***Name***	Type a unique name
	Type	**ICMP**
	Interval	**30** (recommended)
	Timeout	**91** (recommended)
Pool (*Local Traffic -->Pools*)	***Name***	Type a unique name
	Health Monitor	Select the appropriate monitor you created
	Slow Ramp Time	**300**
	Load Balancing Method	Choose a load balancing method. We recommend **Least Connections (Member)**
	Address	Type the IP Address of a server.
	Service Port	Type the appropriate port, such as UDP port **514**, the port on which logging typically occurs. Click **Add**, and then repeat Address and Port for all nodes

2. Log into the BIG-IP system using the command line. Enter the tmsh shell, by typing **tmsh** from the prompt.
3. Create a Remote High Speed Log (HSL) destination:

```
(tmos)# create / sys log-config destination remote-high-speed-log [name] pool-name [specified pool] protocol [udp or tcp]
```

4. If you have a specific log format requirement, create a format-specific log destination, and forward that to the previously-created HSL destination:

```
(tmos)# create / sys log-config destination [splunk|arcsight|remote-high-speed-log] [name] forward-to [HSL name]
```

5. Create a log publisher:

```
(tmos)# create / sys log-config publisher [name] destinations add { [logdestination name] }
```

6. Create the logging profile to tie everything together.
   If you chose to log allowed connections, include the green text (as in step 2 substep l in *To configure the BIG-IP AFM to allow connections from a single trusted network on page 42)*.
   If you set the rule to drop incoming connections, include the text in blue.
   If you chose to log IP intelligence events, include the text in red to add the parameter that sets the log publisher.

```
(tmos)# create / security log profile [name] network add { [name] { filter { log-acl-match-accept enabled
log-acl-match-drop enabled log-acl-match-reject enabled } format { field-list { date_time action drop_reason
protocol src_ip src_port dest_ip dest_port } type field-list } publisher [logpublisher name] } } ip-
intelligence { log-publisher [logpublisher name] }
```

Assigning the logging profile to the virtual server

The final task is to assign the logging profile to the virtual server.

To assign the logging profile to the Skype virtual server

1. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.
2. Click the name of your Skype for Business External Edge or reverse proxy virtual server.
3. From the **Security** menu, choose **Policies**.
4. Next to **Log Profile**, select **Enabled**, then select the Logging profile you created.
5. Click **Update**. The list screen and the updated item are displayed.

This completes the manual configuration.

Revision History

Version	Description	Date
1.0	New guide for Microsoft Skype for Business Server 2015.	07-06-2015

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