



Deploying the BIG-IP LTM with Microsoft Office Communications Server 2007 R2

Table of Contents

Introducing the F5 and Microsoft Office Communications Server 2007 R2 configuration	
Prerequisites and configuration notes	1-1
Product versions and revision history	1-2
Configuration example	1-2
 Configuring the BIG-IP LTM for Microsoft Office Communications Server 2007 R2	
Performing the initial configuration tasks	1-4
Configuring the BIG-IP LTM for the Front End servers	1-7
Configuring the BIG-IP LTM for HTTPS/SSL (444) traffic on the Front End servers	1-8
Creating the Front End health monitor	1-8
Creating the Front End SSL pool	1-9
Creating the Front End profiles	1-11
Creating the Front End virtual server	1-13
Creating a SNAT	1-14
Configuring additional protocols and services for Front End servers	1-17
Synchronizing the BIG-IP configuration if using a redundant system	1-23
 Configuring the BIG-IP LTM for Communicator Web Access	
Importing keys and certificates	2-1
Creating the HTTP health monitor	2-2
Creating the CWA pool	2-2
Creating profiles	2-3
Creating the virtual server	2-7
 Configuring the BIG-IP LTM for Office Communications Server R2 Edge servers	
Configuring the BIG-IP LTM for HTTPS/SSL (444) traffic on the Edge servers	3-2
Creating the Edge server health monitor	3-2
Creating the Edge server HTTPS/SSL pool	3-2
Creating the Edge server TCP profile	3-3
Creating the Edge server virtual server	3-4
Configuring additional protocols and services on the Edge servers	3-6



I

Deploying F5 with Microsoft Office Communications Server R2

- Configuring the BIG-IP LTM for Microsoft Office Communications Server 2007 R2
- Performing the initial configuration tasks
- Configuring the BIG-IP LTM for the Front End servers
- Configuring the BIG-IP LTM for HTTPS/SSL (444) traffic on the Front End servers
- Configuring additional protocols and services for Front End servers

Introducing the F5 and Microsoft Office Communications Server 2007 R2 configuration

Welcome to the Microsoft® Office Communications Server 2007 R2 deployment guide. This guide contains step-by-step procedures for configuring the BIG-IP LTM system with Microsoft Office Communications Server 2007 R2.

This deployment guide is the result of collaboration and interoperability testing between Microsoft and F5 Networks using Microsoft Office Communications Server 2007 R2 and the BIG-IP Local Traffic Manager (LTM). Organizations using the BIG-IP LTM system benefit from mission-critical availability, intelligent traffic management, simple scalability, and enhanced security for Office Communications Server deployments.

For more information on Microsoft Office Communications Server, see <http://office.microsoft.com/en-us/communicationsserver/default.aspx>, or Microsoft's TechNet documentation at <http://technet.microsoft.com/en-us/office/bb267356.aspx>

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

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

Prerequisites and configuration notes

The following are prerequisites for this deployment:

- ◆ The BIG-IP LTM system must be running version v9.0 or later. We highly recommend using version 9.4 or later. Examples shown in this document are from a v10.0 system, but other than minor interface differences are applicable to v9.x systems as well.
- ◆ You must be running Microsoft Office Communications Server 2007 R2. For deployment guidance for Microsoft Live Communications Server 2005 for BIG-IP versions 4.5 and 9.0, and the initial release of Office Communication Server 2007, see <http://www.f5.com/solutions/>.
- ◆ This document is written with the assumption that you are familiar with both the BIG-IP LTM system and the Office Communications Server 2007 R2. For more information on configuring these products, consult the appropriate documentation.

Product versions and revision history

Product and versions tested for this deployment guide:

Product Tested	Version Tested
BIG-IP LTM	v10.0, 10.1 (also applicable to v9.x)
Office Communications Server 2007	R2

Revision history:

Document Version	Description
1.0	New deployment guide
1.1	<ul style="list-style-type: none"> - Added new section on load balancing Communicator Web Access (CWA) - Added Persistence profile to Front End and Edge roles - Added load balancing method guidance for all services - Modified the health monitors to align with Microsoft guidance - Modified TCP timeouts from 1200 to 1800 seconds. - Divided the deployment guide into chapters
1.2	Revised and clarified SIP health monitor guidance for Front End and Director Servers (pages 1-18 and 1-19)

Configuration example

The BIG-IP LTM system can be used to add high availability and traffic direction to an Office Communication Server 2007 Enterprise Pool. Additionally, the BIG-IP LTM system provides required SNAT functionality to enable inter-server communication within the pool.

The following example shows a typical configuration with a BIG-IP LTM system and an Office Communications Server deployment. With multiple Office Communications Servers in a pool there is a need for distributing the incoming session requests among the servers. Figure 1.1 shows a logical configuration diagram.

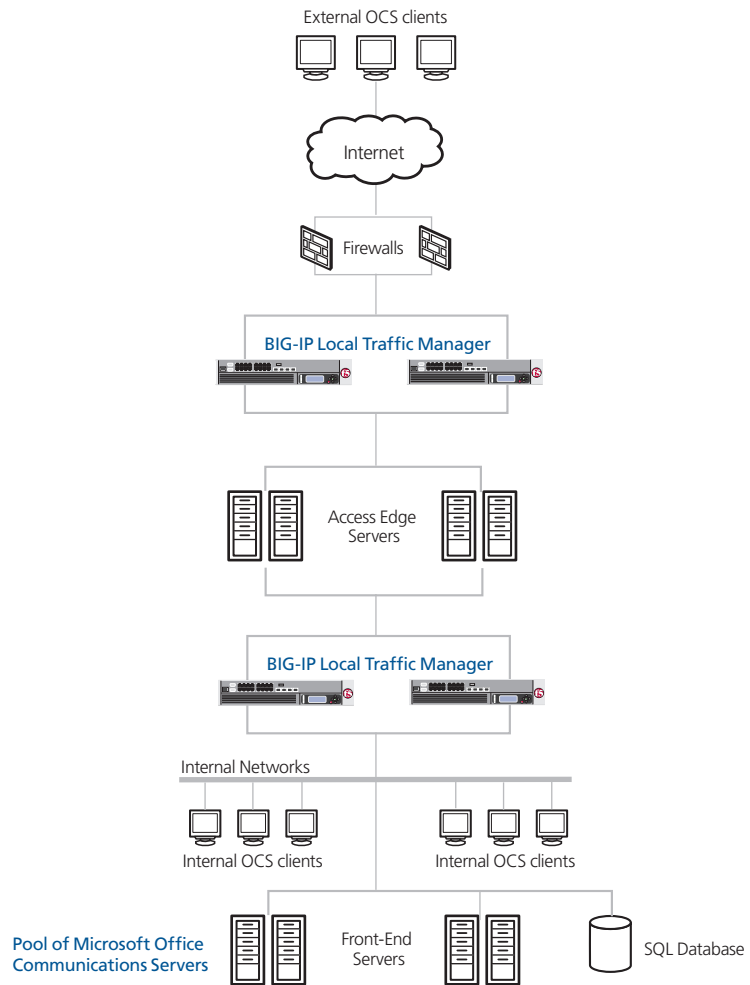


Figure 1.1 BIG-IP LTM and Office Communications Server R2 logical configuration example

Configuring the BIG-IP LTM for Microsoft Office Communications Server 2007 R2

This deployment guide is divided into the following sections:

- *Performing the initial configuration tasks*, on page 1-4
- *Configuring the BIG-IP LTM for the Front End servers*, on page 1-7
- *Configuring the BIG-IP LTM for Communicator Web Access*, on page 2-1
- *Configuring the BIG-IP LTM for Office Communications Server R2 Edge servers*, on page 3-1

We recommend you save your existing BIG-IP configuration before you begin the procedures in this Deployment Guide. For information on backing up or restoring a BIG-IP LTM configuration, refer to the appropriate BIG-IP LTM manual, available on [Ask F5](#).

Performing the initial configuration tasks

In this section, we configure the BIG-IP LTM with a VLAN and Self IP address. Complete these procedures if you do not already configured these objects on the BIG-IP LTM.

Creating a VLAN

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

◆ **One armed configuration**

When the Communicator 2007 R2 clients reside on the same network as the Office Communications Server Front End servers, or you wish to have your BIG-IP LTM virtual servers reside on the same network as your Front End servers, you only need one VLAN. This is also known as a one armed configuration.

◆ **Note**

*In deployments with more than 65,000 simultaneous connections, you need to configure more than one SNAT address on the BIG-IP LTM. See **Creating a SNAT**, on page 14.*

◆ **Routed configuration**

A more common example is when the Communicator 2007 R2 clients and the IP addresses of your BIG-IP LTM virtual servers reside on a different network than the Office Communications Server Front End

servers. In this case, you will need an external VLAN for the incoming clients, and an internal VLAN for the Office Communications Server Front End servers. This is known as a routed configuration.

To create a VLAN

1. On the Main tab, expand **Network**, and then click **VLANs**.
The VLANs screen opens.
2. Click the **Create** button.
The new VLAN screen opens.
3. In the **Name** box, type a unique name for the VLAN. In our example we use **ocs-vlan**.
4. In the **Tag** box, you can optionally type a tag. In our example, we leave this blank, and the BIG-IP LTM automatically assigns a tag.
5. In the Resources section, from the **Available** list, select the interface that will have access to tagged traffic, and add it to the **Untagged** box by clicking the Add (<<) button.
In our example, we select **1.14**.
6. Click the **Finished** button.

The screenshot shows the 'New VLAN...' configuration window in the BIG-IP LTM Configuration utility. The breadcrumb path is 'Network >> VLANs >> New VLAN...'. Under 'General Properties', the 'Name' field contains 'ocs-vlan' and the 'Tag' field is empty. The 'Resources' section has three columns: 'Untagged', 'Available', and 'Tagged'. The 'Available' column lists interfaces 1.13, 1.15, 1.16, 2.1, and 2.2. The 'Untagged' column contains interface 1.14. Navigation arrows (<< and >>) are visible between the columns. Below the resources, the 'Configuration' dropdown is set to 'Basic'. The 'Source Check' checkbox is unchecked, and the 'MTU' field is set to '1500'. At the bottom, there are 'Cancel', 'Repeat', and 'Finished' buttons.

Figure 1.2 Adding a VLAN in the BIG-IP LTM Configuration utility

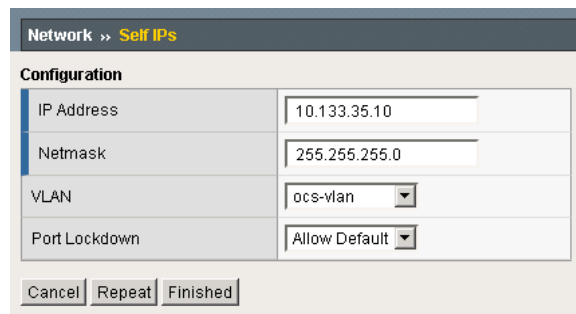
If you are using a routed configuration, you need at least two VLANs on the BIG-IP LTM system. Use the preceding procedure to create each VLAN. Give each VLANs a distinct name (such as **ocs-internal-vlan** and **ocs-external-vlan**), and assign them to the interfaces through which each VLANs traffic should flow.

Creating a self IP

Self IP addresses are the IP addresses owned by the BIG-IP LTM system that you use to access the internal and external VLANs. The next step in this configuration is to create a self IP address for the VLAN we created in the preceding procedure.

To create a self IP address

1. On the Main tab, expand **Network**, and then click **Self IPs**.
The Self IP screen opens.
2. Click the **Create** button.
The new Self IP screen opens.
3. In the **IP Address** box, type a static IP address that resides on the VLAN you created in the preceding procedure. Note that this needs to be on the same network as the Office Communications Server R2 *Servers*. In our example, we use **10.133.35.10**.
4. In the **Netmask** box, type the corresponding subnet mask.
In our example, we use **255.255.255.0**.
5. From the **VLAN** list, select the VLAN you created in *Creating a VLAN*. In our example, we select **ocs-vlan**.
6. Click the **Finished** button.
7. The new self IP address appears in the list.



The screenshot shows the 'Network >> Self IPs' configuration page. It features a 'Configuration' section with four rows: 'IP Address' (10.133.35.10), 'Netmask' (255.255.255.0), 'VLAN' (ocs-vlan), and 'Port Lockdown' (Allow Default). At the bottom are 'Cancel', 'Repeat', and 'Finished' buttons.

Configuration	
IP Address	10.133.35.10
Netmask	255.255.255.0
VLAN	ocs-vlan
Port Lockdown	Allow Default

Buttons: Cancel Repeat Finished

Figure 1.3 Adding a self IP address in the BIG-IP Configuration utility

In a routed configuration, repeat this procedure to add a self IP to each VLAN you created.

This completes the initial configuration section.

Configuring the BIG-IP LTM for the Front End servers

In the following procedures, we configure the BIG-IP LTM for the Office Communications Server R2 Front End servers. A Front End server is an Office Communications Server 2007 R2 server that is a member of an Enterprise Pool that provides the following services: IM Conferencing, Address Book, and Telephony Conferencing.

A Front End Servers in a supported Enterprise Edition configuration can run one or more roles per computer. In the following configurations, we demonstrate how to add support for each role or service to a BIG-IP LTM. Your actual network topology, intended user base, and infrastructure requirements will dictate your deployment scenario and the computers to which you assign each role; you should consult the documentation at [http://technet.microsoft.com/en-us/library/dd425322\(Office.13\).aspx](http://technet.microsoft.com/en-us/library/dd425322(Office.13).aspx) for examples of supported topologies.

Front End servers and the roles they provide can be configured in a variety of topologies. The steps for configuring the BIG-IP LTM are substantially similar for each role, differing only in details such as port, monitor type, and pool members. Rather than repeat the set of instructions for each role, we provide a detailed example for one virtual server and its associated objects and parameters. Following the example, we provide tables of other required and optional roles and services which you may choose to deploy. For each role, you follow the same steps as our examples, modifying indicated parameters or following deployment notes as required for each. Services that share the same port may not also share the same IP address, though services with different ports may. The table indicates the minimum number of unique IP addresses required for the BIG-IP LTM virtual servers.

For the Front End (FE) servers, we provide the following configuration instructions and data:

- *Configuring the BIG-IP LTM for HTTPS/SSL (444) traffic on the Front End servers*, on page 1-8
- *Configuring additional protocols and services for Front End servers*, on page 1-17
- *Synchronizing the BIG-IP configuration if using a redundant system*, on page 1-23

Configuring the BIG-IP LTM for HTTPS/SSL (444) traffic on the Front End servers

Microsoft Office Communications Server R2 servers require HTTPS communication with Front End servers on custom port **444**. Using the BIG-IP LTM to direct traffic to a pool of Front End servers provides load distribution, high availability, and increased scalability.

Creating the Front End health monitor

The first step in configuring the BIG-IP LTM for the Front End servers is to configure a health monitor on the BIG-IP LTM system. We use the HTTPS parent monitor to create this monitor.

◆ **Note**

Microsoft recommends monitoring port 5061 to ascertain the availability of the service, even though the pool uses an alternate service port.

To configure a health monitor

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**. The Monitors screen opens.
2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a name for the Monitor. In our example, we type **ocs-fe-tcp**.
4. From the Type list, select **TCP**. The HTTPS monitor configuration options appear.
5. From the **Configuration** list, select **Advanced**. The advanced configuration options appear.
6. In the **Interval** box, type an number of seconds. This indicates the frequency of the health check. In our example, we type **30**.
7. In the **Timeout** box, type a number of seconds. We recommend at least a 1:3 +1 ratio between the Interval and the Timeout. In our example, we type **91**.
8. In the **Alias Service Port** box, type **5061**.
9. All other configuration settings are optional, configure as applicable for your deployment.
10. Click the **Finished** button (see Figure 1.4).

Local Traffic » Monitors » New Monitor ...

General Properties

Name	ocs-fe-tcp
Type	TCP
Import Settings	tcp

Configuration: Advanced

Interval	30 seconds
Up Interval	Disabled
Time Until Up	0 seconds
Timeout	91 seconds
Manual Resume	<input type="radio"/> Yes <input checked="" type="radio"/> No
Send String	<input style="width: 100%;" type="text"/>
Receive String	<input style="width: 100%;" type="text"/>
Transparent	<input type="radio"/> Yes <input checked="" type="radio"/> No
Alias Address	* All Addresses
Alias Service Port	5061 Other: ▼

Figure 1.4 Configuring the Front End health monitor

Creating the Front End SSL pool

The next step is to create a pool on the BIG-IP LTM system for the Office Communications Server Front End servers. A BIG-IP LTM pool is a set of devices grouped together to receive traffic according to a load balancing method.

To configure the Front End pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**. The Pool screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New Pool screen opens.
3. From the **Configuration** list, select **Advanced**. The advanced configuration options appear.
4. In the **Name** box, type a name for your pool. In our example, we use **ocs-frontend-ssl**.

5. In the **Health Monitors** section, from the **Available** list, select the name of the monitor you created in *Creating the Front End health monitor*, on page 8, and click the Add (<<) button. In our example, we select **ocs-fe-https**.
6. *Optional:* In the **Slow Ramp Time** box, type **300**.
Because we are using the Least Connections load balancing method, we set the Slow Ramp Time in order to ensure that if a pool member becomes available after maintenance or a new member is added, the BIG-IP does not send all new connections to that member (a newly available member always has the least number of connections).
7. In the Resources section, from the **Load Balancing Method** list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network).
For this configuration, we recommend selecting **Least Connections (node)**. In Least Connections mode, the BIG-IP LTM system passes a new connection to the node that has the least number of current connections. Least Connections mode works best in environments where the servers or other equipment you are load balancing have similar capabilities. With Office Communications Server, traffic from servers to clients is roughly the same on each connection.
8. In the New Members section, you add the Office Communications Front End servers to the pool.
 - a) In the **Address** box, type the IP address of the Front End server. In our example, we type **10.133.35.21**.
 - b) In the **Service Port** box, type the service number you want to use for this device. In our example, we type **444**.
 - c) Click the **Add** button to add the member to the list.
 - d) Repeat steps a-c for each device you want to add to the pool. In our example, we repeat these steps for the other two Front End servers (**10.133.35.22** and **10.133.35.23**).
9. Click the **Finished** button (see Figure 1.5).

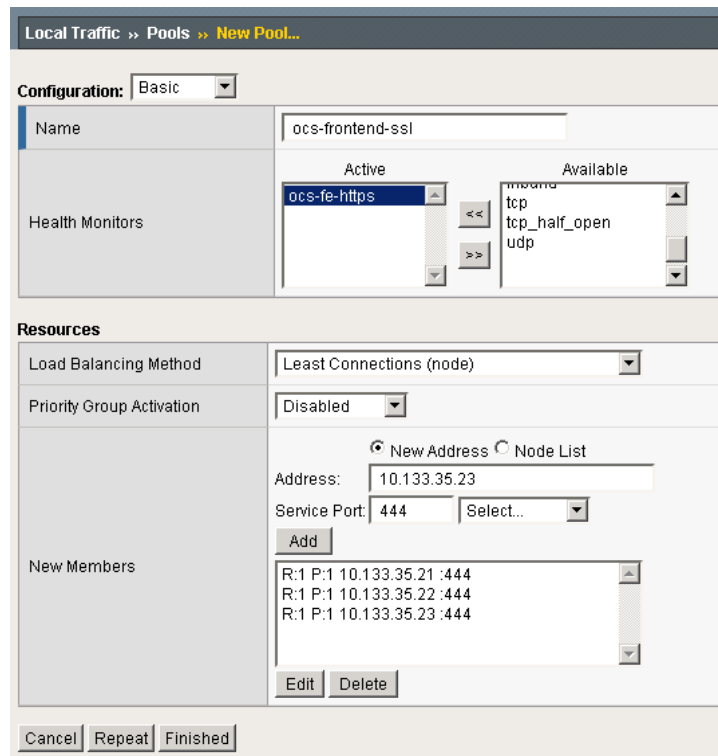


Figure 1.5 Creating the Front End pool

Creating the Front End profiles

The next task is to create the profiles. A profile is an object that contains user-configurable settings for controlling the behavior of a particular type of network traffic. Using profiles enhances your control over managing network traffic, and makes traffic-management tasks easier and more efficient. We create TCP and Persistence profiles.

Creating the TCP profile

The first profile we create is a TCP profile. If most of your end users are on the LAN (which is typical for Front End servers), we recommend using the **tcp-lan-optimized** parent profile.

No matter which TCP profile parent you use, we recommend setting the **Idle Timeout** value to period of time longer than the default (we use an Idle Timeout of 1800 seconds in our example). If a connection is completely idle for this period, the BIG-IP LTM system resets the connection. We set the Idle Timeout value higher than the default setting because it is important to allow connections to remain open and idle for longer time periods, as this is normal behavior of Office Communications Server R2 clients.

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, from the **Protocol** menu, select **TCP**.
3. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **ocs-fe-ssl-lan-opt**.
5. From the **Parent Profile** list, select **tcp-lan-optimized**.
6. In the **Idle Timeout** row, check the **Custom** box. Leave the list set to **Specify**, and in the box, type **1800**.
7. Modify the rest of the settings as applicable for your network. The default settings should suffice for most networks.
8. Click the **Finished** button.

General Properties	
Name	ocs-fe-ssl-lan-opt
Parent Profile	tcp-lan-optimized

Settings		Custom <input type="checkbox"/>
Reset On Timeout	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/>
Time Wait Recycle	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/>
Delayed Acks	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/>
Proxy Maximum Segment	<input type="checkbox"/>	<input type="checkbox"/>
Proxy Options	<input type="checkbox"/>	<input type="checkbox"/>
Proxy Buffer Low	98304 bytes	<input type="checkbox"/>
Proxy Buffer High	131072 bytes	<input type="checkbox"/>
Idle Timeout	Specify... 1800 seconds	<input checked="" type="checkbox"/>
Time Wait	Specify... 2000 milliseconds	<input type="checkbox"/>
Fin Wait	Specify... 5 seconds	<input type="checkbox"/>

Figure 1.6 Creating a new TCP profile

Creating the persistence profile

The next profile we create is a persistence profile. We recommend using Source Address Affinity.

To create the persistence profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, click **Persistence**.
3. Click the **Create** button.

-
4. In the **Name** box, type a name. In our example, we type **ocs-fe-ssl-persist**.
 5. From the **Persistence Type** list, select **Source Address Affinity**.
 6. Configure any of the options as applicable. In our example, we leave the defaults.
 7. Click **Finished**.

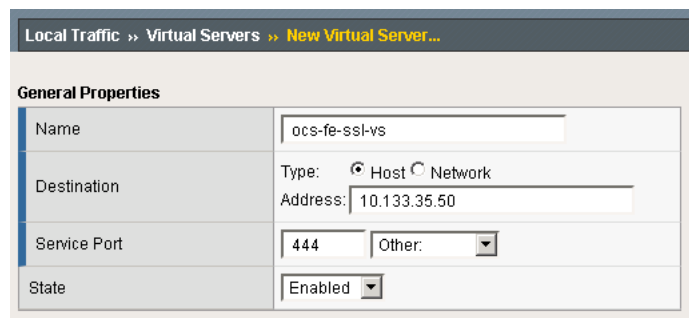
Creating the Front End virtual server

A virtual server with its virtual address is the visible, routable entity through which the Office Communications Servers in a load balancing pool are made available to the client (the IP address to give clients or add to DNS). In the Office Communications Server R2 documentation, this is referred to as the Hardware Load Balancer (HLB) virtual IP (VIP).

The next step in this configuration is to define a virtual server that references the profile and pool you created.

To create the Front End SSL virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New Virtual Server screen opens.
3. In the **Name** box, type a name for this virtual server. In our example, we type **ocs-fe-ssl-vs**.
4. In the Destination section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we use **10.133.35.50**.
6. In the Service Port box, type **444**.



The screenshot shows the 'New Virtual Server' configuration window. The breadcrumb path is 'Local Traffic >> Virtual Servers >> New Virtual Server...'. The 'General Properties' section is expanded, showing the following fields:

Name	ocs-fe-ssl-vs
Destination	Type: <input checked="" type="radio"/> Host <input type="radio"/> Network Address: 10.133.35.50
Service Port	444 Other: <input type="text"/>
State	Enabled <input type="text"/>

Figure 1.7 The General Properties of the Front End SSL virtual server

7. From the **Configuration** list, select **Advanced**.

8. From the **Protocol Profile (Client)** list, select the name of the profile you created in *Creating the Front End profiles*, on page 1-11. In our example, we select **ocs-fe-ssl**.
9. In the Resources section, from the **Default Pool** list, select the pool you created in *Creating the Front End SSL pool*, on page 1-9. In our example, we select **ocs-frontend-ssl**.
10. From the Default Persistence Profile list, select the profile you created in *Creating the persistence profile*, on page 1-12. In our example, we select **ocs-fe-ssl-persist**.
11. Click the **Finished** button.

Up Down	
Default Pool	ocs-frontend-ssl
Default Persistence Profile	ocs-fe-ssl-persist
Fallback Persistence Profile	None
Cancel Repeat Finished	

Figure 1.8 Resources section of the Front End SSL virtual server

Creating a SNAT

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

A default SNAT is appropriate for most deployments. If more than 65,000 simultaneous users are connecting to the Office Communications Server deployment, see *Configuring a SNAT for large Office Communications Server deployments*, on page 1-15.

Use the procedure most applicable for your deployment.

To create a default SNAT for less than 65,000 concurrent users

Use this procedure if your Office Communications Server deployment has fewer than 65,000 simultaneous users.

To create a default SNAT

1. On the Main tab, expand **Local Traffic**, and then click **SNATs**. The SNATs screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New SNAT screen opens.

3. In the **Name** box, type a name for this SNAT. In our example, we type **ocs-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 Office Communications Server R2 devices reside, and click the Add (<<) button. In our example, we select **ocs-vlan**.
7. Click the **Finished** button.

Figure 1.9 Creating the default SNAT

Configuring a SNAT for large Office Communications Server deployments

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

◆ Important

This procedure is only necessary for large deployments. If your Office Communications Server R2 deployment has less than 65,000 simultaneous users, you do not need to create a SNAT pool. Use the previous procedure.

To create a SNAT pool for large deployments

1. On the Main tab, expand Local Traffic, and then click SNATs. The SNATs screen opens.
2. On the Menu bar, click **SNAT Pool List**.

3. In the upper right portion of the screen, click the **Create** button. The New SNAT Pool screen opens.
4. In the **Name** box, type a name for this SNAT Pool. In our example, we type **ocs-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. In our example, we type **10.133.35.110**.

Repeat this step for each additional address needed. At least one address should be added for each 65,000 anticipated concurrent connections (the number of connection generally corresponds to the number of clients). In our example, we add **10.133.35.111**.

6. Click the **Finished** button.




Figure 1.10 Creating the SNAT pool for large deployments

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**. The SNATs screen opens.
8. In the upper right portion of the screen, click the **Create** button. The New SNAT screen opens.
9. In the **Name** box, type a name for this SNAT. In our example, we type **ocs-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 **ocs-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 Office Communications Server R2 devices reside, and click the Add (<<) button. In our example, we select **ocs-vlan**.
 14. Click the **Finished** button.

This completes the HTTPS/SSL portion of the Front End server portion of the configuration. Continue with the following section.

Configuring additional protocols and services for Front End servers

There are a number of required and optional services and protocols for Front End servers. This section contains four tables which list these services and protocols. Review each table, and for the services and protocols applicable to your configuration, repeat the procedures in *Configuring the BIG-IP LTM for the Front End server additional protocols and services*, on page 1-19, using the port, protocol, monitor type and other information from the tables, as appropriate. Be sure to define nodes and pools that are correct according to which servers you have elected to install each service.

Services that share the same port may not share the same IP address, though services with different ports may. The table indicates the minimum number of unique IP addresses required for the virtual servers.

◆ Note

We recommend the Least Connections (Node) load balancing method for all pools on the BIG-IP LTM.

How to use the following tables

For each of the line items in the following tables, you must create a health monitor, pool, TCP profile, and virtual server on the BIG-IP LTM. Each table uses the same virtual server IP address on the BIG-IP LTM system, though the port numbers are different. Use a unique name that corresponds to the service for each of the new BIG-IP LTM objects you configure. This helps avoid confusion. For example, you might use **ocs-fe-sip-virtual** or **ocs-fe-https-virtual**, as easily identifiable names for two virtual servers.

◆ Role:Pool Members

This describes the Microsoft Office Communications Server R2 Role, and the devices that are used for the BIG-IP LTM configuration objects.

◆ Service

This is the name of the service or protocol used by the Role.

◆ VIP TCP Port

The VIP TCP Port is the Service Port you enter when configuring the BIG-IP LTM virtual server.

◆ **Pool TCP Port**

The Pool TCP Port is the Service Port you enter when configuring the BIG-IP LTM Pool.

◆ **Monitor Port (Monitor Type)**

The Monitor Port (Monitor Type) column contains both the Monitor Port, which you enter when configuring the Alias Service Port for the monitor, and the Monitor Type, which you select when configuring the Monitor Type.

Virtual Server IP Address 1: OCS FE Enterprise Pool

Role:Pool Members	Service	Monitor Port (Monitor Type)	Pool TCP Port	VIP TCP Port
OCS R2 Enterprise Pool, Required Services: OCS R2 Front End Enterprise Pool	SIP	5060 (SIP) or 5061 (TCP)*	5061	5061
	HTTPS	443 (HTTPS)	443	443
	HTTPS-444	444 (HTTPS: 5061*)	444	444
	DCOM	135 (TCP: 5061*)	135	135
	APPSHARE	5065 (TCP)	5065	5065
	QOE-AGENT	5069 (TCP)	5069	5069

See the **Note in Creating the Front End health monitors, on page 1-19 for an explanation.*

***Microsoft recommends monitoring port 5061 to ascertain the availability of these services, even though the pools use alternate service ports.*

Virtual Server IP Address 2: OCS Director

Role:Pool Members	Service	Monitor Port (Monitor Type)	Pool TCP Port	VIP TCP Port
OCS Director: OCS Director Servers	SIP-TCP	5060 (SIP)	5060	5060
	SIP-TLS	5060 (SIP) or 5061 (TCP)*	5061	5061

See the **Note in Creating the Front End health monitors, on page 1-19 for an explanation.*

◆ **Note**

The Director servers will almost always be different from those servers running other Front End services.

Virtual Server IP Address 3: Various Roles

Role:Pool Members	Service	Monitor Port (Monitor Type)	Pool TCP Port	VIP TCP Port
OCS Response Group Service: OCS FE Servers with the Response Group Service	SIP	5071 (TCP)	5071	5071
OCS Dial-in Conferencing Support: OCS FE Servers with the Conferencing Attendant	SIP-5072	5072 (TCP)	5072	5072
	SIP-5073	5073 (TCP)	5073	5073
OCS Outside Voice Control: OCS FE Servers with Outside Voice Control	SIP-5074	5074 (TCP)	5074	5074

Configuring the BIG-IP LTM for the Front End server additional protocols and services

Use the following procedures as a template for configuring the Front End server protocols and services applicable to your configuration, as described in the tables above.

Creating the Front End health monitors

The first step is to configure the health monitor. This procedure uses entries from the **Monitor Port (Monitor Type)** column in the tables above.

◆ Note

*For Front End SIP health monitors where the TCP port for the pool members and virtual server is **5061**, administrators may elect to health check using a SIP monitor on port **5060** (the SIP-TCP service) rather than using a simple TCP monitor on port **5061** (the SIP-TLS service) in order to avoid creating TLS error entries on the Front End Server application logs. While this is a simple change, administrators should be aware that port **5060** is not enabled by default, and the health status of port **5060** may not necessarily accurately reflect the health of the **5061** TLS listener, even though they are normally in the same state. We recommend continuing to use a simple TCP monitor on port **5061** (the SIP-TLS service).*

To configure a health monitor for the Front End servers

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**. The Monitors screen opens.
2. Click the **Create** button. The New Monitor screen opens.

3. In the **Name** box, type a unique name for this Monitor.
We recommend using *FE* or *Front-end* and the *protocol* in the name to avoid confusion. For example, **ocs-fe-sip**.
4. From the **Type** list, select the monitor type found in the **Monitor Port (Monitor Type)** column in parenthesis. For example, if the column contains **5061 (SIP)**, select **SIP** from the **Type** list.
The Monitor configuration options appear.
5. From the **Configuration** list, select **Advanced**.
6. In the Configuration section, in the **Interval** and **Timeout** boxes, type an Interval and Timeout. We recommend at least a 1:3 +1 ratio between the interval and the timeout. In our example, we use a Interval of **30** and a Timeout of **91**.
7. In the **Alias Service Port** box, type the appropriate port found in the **Monitor Port (Monitor Type)** column. For example, if the column contains **5061 (SIP)**, type **5061** in the **Alias Service Port** box.
8. All other configuration settings are optional, configure as applicable for your deployment.
9. Click the **Finished** button.

Creating the Front End pools

The next step is to create the pools on the BIG-IP LTM system. This procedure uses entries from the **Pool TCP Port** column in the tables above.

Creating the Front End server pools

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
The Pool screen opens.
2. Click the **Create** button. The New Pool screen opens.
3. From the **Configuration** list, select **Advanced**.
4. In the **Name** box, type a unique name for this Pool.
We recommend using *FE* or *Front-end* and the *protocol* in the name to avoid confusion. For example, **ocs-fe-sip**.
5. In the Health Monitors section, from the **Available** list, select the name of the monitor you created in *Creating the Front End health monitors*, on page 19, and click the Add (<<) button. Be sure to use the monitor that is associated with the same service or protocol as this pool.
6. *Optional:* In the **Slow Ramp Time** box, type **300**.
Because we are using the Least Connections load balancing method, we set the Slow Ramp Time in order to ensure that if a pool member becomes available after maintenance or a new member is added, the BIG-IP does not send all new connections to that member (a newly available member always has the least number of connections).

-
7. In the Resources section, from the **Load Balancing Method** list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network). We recommend selecting **Least Connections (node)** for all pools in this configuration.
 8. In the New Members section, you add the Office Communications Front End Servers to the pool.
 - a) In the **Address** box, type the IP address of the Office Communications Front End Server.
 - b) In the **Service Port** box, type the service number from the **Pool TCP Port** column in the table above. For example, if you are configuring the SIP pool, use port **5061**.
 - c) Click the **Add** button to add the member to the list.
 - d) Repeat steps a-c for each Office Communications Front End Server you want to add to the pool.
 9. Click the **Finished** button.

Creating the Front End profiles

The next step is to create the TCP and Persistence profiles.

For the TCP profile, if most of your end users are on the LAN (which is typical for Front End servers), we recommend using the **tcp-lan-optimized** parent profile.

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, from the **Protocol** menu, select **TCP**.
3. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a unique name for this profile. We recommend using *FE* or *Front-end*, the *protocol*, and the type of TCP profile in the name to avoid confusion. For example, **ocs-fe-sip-lan-opt**.
5. From the **Parent Profile** list, select **tcp-lan-optimized**.
6. In the **Idle Timeout** row, check the **Custom** box. Leave the list set to **Specify**, and in the box, type **1800**.
7. Modify the rest of the settings as applicable for your network. The default settings should suffice for most networks.
8. Click the **Finished** button.

To create the persistence profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.

2. On the Menu bar, click **Persistence**.
3. Click the **Create** button.
4. In the **Name** box, type a unique name for this profile.
5. From **Persistence Type** list, select **Source Address Affinity**.
6. Configure any of the options as applicable. In our example, we leave the defaults.
7. Click **Finished**.

Creating the Front End virtual servers

The final step in this section is to define a virtual server that references the profile and pool you created. This procedure uses entries from the **VIP TCP Port** column in the tables above.

To create the Front End virtual servers

1. On the Main tab, expand **Local Traffic**, and click **Virtual Servers**.
2. Click the **Create** button. The New Virtual Server screen opens.
3. In the **Name** box, type a unique name for this Virtual Server. We recommend using *FE* or *Front-end* and the *protocol* in the name to avoid confusion. For example, **ocs-fe-sip-vs**.
4. In the Destination section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we use **10.133.36.51**.
6. In the **Service Port** box, type the service number from the **VIP TCP Port** column in the table above. For example, if you are configuring the SIP virtual server, use port **5061**.
7. From the **Configuration** list, select **Advanced**.
8. From the **Protocol Profile (Client)** list, select the profile you created in *Creating the Front End profiles*, on page 1-21.
9. In the Resources section, from the **Default Pool** list, select the pool you created in *Creating the Front End pools*, on page 1-20.
10. From the **Default Persistence Profile** list, select the profile you created in *Creating the Front End profiles*, on page 1-21.
11. Configure any other settings as appropriate for your configuration.
12. Click the **Finished** button.
Repeat this section for each of the Front End protocols and services in your configuration.

Synchronizing the BIG-IP configuration if using a redundant system

When you have completed the configuration of your virtual servers and related objects, and if you are using a redundant BIG-IP configuration, the final step is to synchronize the configuration to the peer BIG-IP device.

The method of synchronizing the BIG-IP configuration depends on your version, see the appropriate BIG-IP LTM manual, available on Ask F5 (https://support.f5.com/kb/en-us/products/big-ip_ltm.html)

◆ Important

If you have a redundant BIG-IP configuration (active-active or active-standby), you must also perform the first two procedures (Creating a VLAN and Creating a self IP) on both devices. The rest of the procedures only need to be performed on one BIG-IP device. The first two procedures are not included in the items that are synchronized between the BIG-IP devices.

In a redundant configuration, you also need to configure a Floating Self IP address for the VLAN on both devices. To create this Floating Self IP address, follow the procedure Creating a self IP, on page 5, but check the Floating IP box. On the redundant device, create a Floating Self IP address using the same IP address as the original device, and check the Floating IP box.



2

Deploying F5 with Microsoft Office Communications Server R2 CWA

- Configuring the BIG-IP LTM for Communicator Web Access

Configuring the BIG-IP LTM for Communicator Web Access

In this chapter, we configure the BIG-IP LTM for the Communicator Web Access component of Office Communications Server R2. Microsoft Office Communicator Web Access enables you to provide Office Communications Server services – such as instant messaging (IM), presence, audio conferencing, and desktop sharing – to users who do not use Office Communicator.

See technet.microsoft.com/en-us/library/dd441196%28office.13%29.aspx for Microsoft guidance, which we have incorporated in the following procedures.

Importing keys and certificates

If you are using the BIG-IP LTM system for offloading SSL from the CWA devices, you must install a SSL certificate and key on the BIG-IP LTM system. For this Deployment Guide, we assume that you already have obtained an SSL certificate, but it is not yet installed on the BIG-IP LTM system. For information on generating certificates, or using the BIG-IP LTM system to generate a request for a new certificate and key from a certificate authority, see the Managing SSL Traffic chapter in the *Configuration Guide for Local Traffic Management*.

Once you have obtained a certificate, you can import this certificate into the BIG-IP LTM system using the Configuration utility. You can use the Import SSL Certificates and Keys screen only when the certificate you are importing is in Privacy Enhanced Mail (PEM) format.

Important

If you are not using the BIG-IP LTM system for offloading SSL, you do not need to perform this procedure.

To import a key or certificate

1. On the Main tab, expand Local Traffic.
2. Click **SSL Certificates**.
This displays the list of existing certificates.
3. In the upper right corner of the screen, click **Import**.
4. From the **Import** list, select the type of import (**Key** or **Certificate**).
5. Select the import method (text or file).
6. Type the name of the key or certificate.
7. Click **Import**.
8. If you imported the certificate, repeat this procedure for the key.

Creating the HTTP health monitor

The first step in this configuration is to set up an HTTP health monitor. This procedure is optional, but very strongly recommended. For this configuration, we use an HTTP monitor, which checks nodes (IP address and port combinations), and can be configured to use Send and Receive strings in an attempt to retrieve explicit content from nodes, as we show in the following example.

◆ Tip

Although we strongly recommend a health monitor, it does not have to be an HTTP monitor. You can also configure multiple health monitors, such as configuring a basic TCP monitor in addition to the HTTP monitor.

To configure the health monitor

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**.
2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a name for the Monitor.
In our example, we type **ocs-cwa-monitor**.
4. From the **Type** list, select **http**.
5. In the Configuration section, in the **Interval** and **Timeout** boxes, type an Interval and Timeout. We recommend at least a 1:3 +1 ratio between the interval and the timeout (for example, the default setting has an interval of **5** and an timeout of **16**). In our example, we use a **Interval** of **30** and a **Timeout** of **91**.
6. In the **Send String** box, you can add an optional Send String specific to the device or application being checked.
7. In the **Receive String** box, type what you expect the server to return as a result of the Send String.
8. Click the **Finished** button.
The new monitor is added to the Monitor list.

Creating the CWA pool

The next step is to create a pool on the BIG-IP LTM system for CWA. A pool is a set of devices grouped together to receive traffic according to a load balancing method.

To create the pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
The Pool screen opens.
2. In the upper right portion of the screen, click the **Create** button.
The New Pool screen opens.
3. From the Configuration list, select **Advanced**.

-
4. In the **Name** box, enter a name for your pool.
In our example, we use **ocs-cwa-pool**.
 5. In the **Health Monitors** section, select the name of the monitor you created in the *Creating the HTTP health monitor* section, and click the Add (<<) button. In our example, we select **ocs-cwa-monitor**.
 6. *Optional:* In the **Slow Ramp Time** box, type **300**.
Because we are using the Least Connections load balancing method, we set the Slow Ramp Time in order to ensure that if a pool member becomes available after maintenance or a new member is added, the BIG-IP does not send all new connections to that member (a newly available member always has the least number of connections).
 7. In the Resources section, from the **Load Balancing Method** list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network).
In our example, we select **Least Connections (node)**.
 8. For this pool, we leave the Priority Group Activation **Disabled**.
 9. In the New Members section, make sure the **New Address** option button is selected.
 10. In the **Address** box, add the first server to the pool. In our example, we type **10.133.13.41**.
 11. In the **Service Port** box, type **80** or select **HTTP** from the list.
 12. Click the **Add** button to add the member to the list.
 13. Repeat steps 10-12 for each server you want to add to the pool.
 14. Click the **Finished** button.

Creating profiles

The next task is to create the profiles.

Creating an HTTP profile

The first new profile we create is an HTTP profile. The HTTP profile contains numerous configuration options for how the BIG-IP LTM system handles HTTP traffic. For deployments where the majority of users accessing the CWA devices are connecting across a WAN, F5 recommends enabling compression and caching on the BIG-IP LTM by using the **http-wan-optimized-compression-caching**. This profile uses specific compression and caching (among other) settings to optimize traffic over the WAN.

To create a new HTTP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
The HTTP Profiles screen opens.

2. Click the **Create** button. The New HTTP Profile screen opens.
3. In the **Name** box, type a name for this profile. In our example, we type **ocs-cwa-http-opt**.
4. From the **Parent Profile** list, select **http-wan-optimized-compression-caching**.
5. *Optional:* If you using the BIG-IP LTM to offload SSL, in the Settings section, check the Custom box for **Redirect Rewrite**, and from the **Redirect Rewrite** list, select **Match**.
6. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating a TCP profile

The next profiles we create are the TCP profiles. If most of the CWA users are accessing the devices via a Local Area Network, we recommend using the **tcp-lan-optimized** (for server-side TCP connections) parent profile. If the majority of the users are accessing the system from remote or home offices, we recommend using an additional TCP profile, called **tcp-wan-optimized** (for client side TCP connections). In our example, we leave these profiles at their default levels; you can configure any of the options as applicable for your network.

For the internal forwarding virtual servers, we recommend creating an additional TCP LAN profile.

Creating the LAN optimized TCP profile

First we configure the LAN optimized profile. If you do not want to use this optimized profile, you can choose the default TCP parent profile.

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Protocol** menu, click **tcp**.
3. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **ocs-cwa-tcp-lan**.
5. From the **Parent Profile** list, select **tcp-lan-optimized**.
6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating the WAN optimized TCP profile

Now we configure the WAN optimized profile. Remember, if most of the users are accessing the system over the LAN or other low latency links, you do not need to create this profile.

To create a new TCP WAN optimized profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Protocol** menu, click **tcp**.
3. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **ocs-cwa-tcp-wan**.
5. From the **Parent Profile** list, select **tcp-wan-optimized**.
6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating the persistence profile

Next, we create the persistence profile. For CWA, we recommend using cookie persistence (HTTP cookie insert).

To create a new cookie persistence profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, click **Persistence**.
3. In the upper right portion of the screen, click the **Create** button. The New Persistence Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **ocs-cwa-cookie**.
5. From the **Persistence Type** list, select **Cookie**. The configuration options for cookie persistence appear.
6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating a Client SSL profile

If you are using the BIG-IP LTM system to offload SSL, you must create an Client SSL profile. This profile contains the SSL certificate and Key information for offloading the SSL traffic.

To create a new Client SSL profile based on the default profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **SSL** menu, select **Client**. The Client SSL Profiles screen opens.
3. In the upper right portion of the screen, click the **Create** button. The New Client SSL Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **ocs-cwa-clientSSL**.
5. In the Configuration section, click a check in the **Certificate** and **Key Custom** boxes.
6. From the **Certificate** list, select the name of the Certificate you imported in the *Importing keys and certificates* section.
7. From the **Key** list, select the key you imported in the *Importing keys and certificates* section.
8. Click the **Finished** button.

For more information on SSL certificates, or creating or modifying profiles, see the BIG-IP documentation.

Creating a OneConnect Profile

The final profile we create is a OneConnect™ profile. OneConnect improves performance by aggregating multiple client requests into a server-side connection pool, enabling client requests to reuse server-side connections. For more information on OneConnect, see SOL7208 (<https://support.f5.com/kb/en-us/solutions/public/7000/200/sol7208.html>) on Ask F5.

To create a OneConnect Profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Other** menu, select **OneConnect**. The OneConnect profile screen opens.
3. In the upper right portion of the screen, click the **Create** button. The New OneConnect Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **ocs-cwa-oneconnect**.
5. Modify any of the settings as applicable for your configuration. In our example, we leave the settings at their default levels.
6. Click the **Finished** button.

Creating the virtual server

Next, we create the HTTPS virtual server.

To create the HTTPS virtual server

1. On the Main tab, expand **Local Traffic**, click **Virtual Servers**, and then click the **Create** button. The New Virtual Server screen opens.
2. In the **Name** box, type a name for this virtual server. In our example, we type **ocs-cwa-https-vs**.
3. In the **Destination** section, select the **Host** option button.
4. In the **Address** box, type the IP address of this virtual server. In our example, we use **10.133.100.82**.
5. In the **Service Port** box, type **443**, or select **HTTPS** from the list.
6. From the Configuration list, select **Advanced**.
7. From the **Protocol Profile (Client)** list, select the name of the profile you created in the *Creating the WAN optimized TCP profile* section. In our example, we select **ocs-cwa-tcp-wan**. This is optional.
8. From the **Protocol Profile (Server)** list, select the name of the profile you created in the *Creating the LAN optimized TCP profile* section. In our example, we select **ocs-cwa-tcp-lan**.
9. From the **OneConnect Profile** list, select the name of the profile you created in the *Creating a OneConnect Profile* section. In our example, we select **ocs-cwa-oneconnect**.
10. From the **HTTP Profile** list, select the name of the profile you created in the *Creating an HTTP profile* section. In our example, we select **ocs-cwa-http-opt**.
11. From the **SSL Profile (Client)** list, select the SSL profile you created in the *Creating a Client SSL profile* section. In our example, we select **ocs-cwa-clientssl**.
12. In the Resources section, from the **Default Pool** list, select the pool you created in the *Creating the CWA pool* section. In our example, we select **ocs-cwa-web-pool**.
13. From the **Default Persistence Profile** list, select the persistence profile you created in the *Creating the persistence profile* section. In our example, we select **ocs-cwa-cookie**.
14. Click the **Finished** button.

This completes the BIG-IP LTM configuration for Communicator Web Access.



3

Deploying F5 with Microsoft Office Communications Server R2 Edge servers

- Configuring the BIG-IP LTM for HTTPS/SSL (444) traffic on the Edge servers
- Configuring additional protocols and services on the Edge servers

Configuring the BIG-IP LTM for Office Communications Server R2 Edge servers

In this chapter, we configure the Edge servers. An Edge server is an Office Communications Server 2007 R2 server in the perimeter network that provides connectivity for external users and public IM connections. Employees traveling, working from home, or in remote offices, use the Edge servers to remotely access the service.

◆ Important

It is possible to deploy Office Communications Server 2007 R2 without using the Edge Servers or services (for example, in an internal only deployment). If your configuration does not include Edge servers, you do not need to complete this section.

As with the Front End Servers, required and optional roles and services may be installed in a variety of topologies, but the configuration for each on BIG-IP LTM is substantially the same.

We provide two specific examples for the Edge servers, and then a section of tables at the end, similar to the Front End servers, with the relevant configuration information:

- *Configuring the BIG-IP LTM for HTTPS/SSL (444) traffic on the Edge servers, on page 3-2*
- *Configuring additional protocols and services on the Edge servers, on page 3-6*

Configuring the BIG-IP LTM for HTTPS/SSL (444) traffic on the Edge servers

The first task in this configuration is to configure the BIG-IP LTM for Edge server HTTPS traffic on port **444**. As previously noted, Microsoft Office Communications Server clients require HTTPS communication with servers on this custom port.

Creating the Edge server health monitor

The first step in configuring the BIG-IP LTM for the Edge servers is to configure a health monitor on the BIG-IP LTM system. We use the HTTPS parent monitor to create this monitor.

To configure a health monitor

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**.
The Monitors screen opens.
2. Click the **Create** button.
The New Monitor screen opens.
3. In the **Name** box, type a name for the Monitor.
In our example, we type **ocs-edge-https**.
4. From the **Type** list, select **HTTPS**.
The HTTPS Monitor configuration options appear.
5. From the **Configuration** list, select **Advanced**.
The advanced configuration options appear.
6. In the Configuration section, in the **Interval** and **Timeout** boxes, type an Interval and Timeout. We recommend at least a 1:3 +1 ratio between the interval and the timeout. In our example, we use a Interval of **30** and a Timeout of **91**.
7. In the **Alias Service Port** box, type **444**.
8. Click the **Finished** button.

Creating the Edge server HTTPS/SSL pool

The next step is to create an Edge server pool on the BIG-IP LTM system.

To create the Edge SSL pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
The Pool screen opens.
2. Click the **Create** button. The New Pool screen opens.
3. From the **Configuration** list, select **Advanced**.

-
4. In the **Name** box, type a name for your pool.
In our example, we use **ocs-edge-ssl**.
 5. In the Health Monitors section, select the name of the monitor you created in *Creating the Edge server health monitor*, on page 3-2, and click the Add (<<) button. In our example, we select **ocs-edge-https**.
 6. *Optional:* In the **Slow Ramp Time** box, type **300**.
Because we are using the Least Connections load balancing method, we set the Slow Ramp Time in order to ensure that if a pool member becomes available after maintenance or a new member is added, the BIG-IP does not send all new connections to that member (a newly available member always has the least number of connections).
 7. In the Resources section, from the **Load Balancing Method** list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network).

For this configuration, we recommend selecting **Least Connections (node)**. In Least Connections mode, the BIG-IP LTM system passes a new connection to the node that has the least number of current connections. Least Connections mode works best in environments where the servers or other equipment you are load balancing have similar capabilities. Using Office Communications Server, traffic from servers to clients is roughly the same on each connection.

8. In the New Members section, you add the Office Communications Servers to the pool.
 - a) In the **Address** box, type the IP address of one of the Office Communications Edge Servers.
In our example, we type **10.133.36.22**.
 - b) In the **Service Port** box, type **444**.
 - c) Click the **Add** button to add the member to the list.
 - d) Repeat steps a-c for each Office Communications Edge Server you want to add to the pool. In our example, we repeat these steps twice for the other two Office Communications Edge Servers (**10.133.36.23** and **10.133.36.24**).
9. Click the **Finished** button.
10. Repeat this procedure for any other Edge services. Be sure to create a unique name for each pool, and use the appropriate IP addresses and Service Port.

Creating the Edge server TCP profile

The next step is to create a TCP profile. For the Edge servers, most of your end users are likely connecting over the wide area network (WAN), so we recommend using the **tcp-wan-optimized** parent profile.

No matter which TCP profile parent you use, we recommend setting the **Idle Timeout** value to period of time longer than the default (1800 seconds in our example). If a connection is completely idle for this period, the BIG-IP LTM system resets the connection. We set the Idle Timeout value higher than the default setting because it is important to allow connections to remain open and idle for longer time periods, as this is normal behavior of Office Communications Server R2 clients.

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, from the **Protocol** menu, select **TCP**.
3. In the upper right portion of the screen, click the Create button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **ocs-edge-ssl-wan-opt**.
5. From the **Parent Profile** list, select **tcp-wan-optimized**.
6. In the **Idle Timeout** row, check the **Custom** box. Leave the list set to **Specify**, and in the box, type **1800**.
7. Modify the rest of the settings as applicable for your network. The default settings should suffice for most networks.
8. Click the **Finished** button.

General Properties		
Name	ocs-edge-ssl-wan	
Parent Profile	tcp-wan-optimized	

Settings			Custom <input checked="" type="checkbox"/>
Reset On Timeout	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/>	
Time Wait Recycle	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/>	
Delayed Acks	<input checked="" type="checkbox"/> Enabled	<input type="checkbox"/>	
Proxy Maximum Segment	<input type="checkbox"/>	<input type="checkbox"/>	
Proxy Options	<input type="checkbox"/>	<input type="checkbox"/>	
Proxy Buffer Low	131072 bytes	<input type="checkbox"/>	
Proxy Buffer High	131072 bytes	<input type="checkbox"/>	
Idle Timeout	Specify... 1800 seconds	<input checked="" type="checkbox"/>	
Time Wait	Specify... 2000 milliseconds	<input type="checkbox"/>	

Figure 3.1 Configuring the WAN optimized TCP profile

Creating the Edge server virtual server

The next step in this configuration is to define a virtual server that references the profile and pool you just created.

To create the Edge HTTPS/SSL virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New Virtual Server screen opens.
3. In the **Name** box, type a name for this virtual server. In our example, we type **ocs-edge-ssl-vs**.
4. In the Destination section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we use **10.133.36.51**.
6. In the **Service Port** box, type **444**.
7. From the **Configuration** list, select **Advanced**.
8. From the Protocol Profile (Client) list, select the name of the profile you created in *Creating the Edge server TCP profile*, on page 3-3. In our example, we select **ocs-edge-ssl**.
9. In the Resources section, from the **Default Pool** list, select the name of the pool you created in *Creating the Edge server HTTPS/SSL pool*, on page 3-2. In our example, we select **ocs-edge-ssl**.
10. Click the **Finished** button.

This completes the BIG-IP LTM configuration for Edge server HTTPS/SSL traffic on port 444. Continue with the following section.

Configuring additional protocols and services on the Edge servers

For each of the required and optional services listed in the following tables for Edge Servers, you configure a health monitor, pool, TCP profile, and virtual server, using the procedures in *Configuring the BIG-IP LTM for the Edge server additional protocols and services*, on page 3-7. Be sure to define nodes and pools that are correct according to which servers you have elected to install each service.

If a protocol or Role is not listed in the following tables, it is not one that Microsoft supports through a load balanced, NAT (network address translation), or SNAT (source network address translation) configuration.

◆ Note

We recommend the Least Connections (Node) load balancing method for all pools on the BIG-IP LTM.

How to use the following tables

As with the Front End servers, for each of the line items in the following tables, you must create a health monitor, pool, TCP profile, and virtual server on the BIG-IP LTM. Each table uses the same virtual server IP address on the BIG-IP LTM system, though the Port numbers are different. Use a unique name that corresponds to the service for each of the new BIG-IP LTM objects you configure. This helps avoid confusion. For example, you might use **ocs-edge-sip-virtual** or **ocs-edge-https-virtual**, as easily identifiable names for two virtual servers.

As with the Front End server virtual servers, services that share the same port may not also share the same IP address, though services with different ports may. The tables indicate the minimum number of unique IP addresses.

The following describe the columns in the following table:

◆ **Role:Pool Members**

This describes the Microsoft Office Communications Server R2 Role, and the devices that are used for the BIG-IP LTM configuration objects.

◆ **Service**

This is the name of the service or protocol used by the role.

◆ **VIP TCP Port**

The VIP TCP Port is the Service Port you enter when configuring the BIG-IP LTM virtual server.

◆ **Pool TCP Port**

The Pool TCP Port is the Service Port you enter when configuring the BIG-IP LTM Pool.

◆ **Monitor Port (Monitor Type)**

The Monitor Port (Monitor Type) column contains both the Monitor Port, which you enter when configuring the Alias Service Port for the monitor, and the Monitor Type, which you select when configuring the Monitor Type.

Edge Virtual Server IP Address 1: OCS Edge Conferencing servers

Role:Pool Members	Service	Monitor Port (Monitor Type)	Pool TCP Port	VIP TCP Port
OCS R2 Edge Servers: OCS R2 Edge Conferencing Servers	HTTPS-internal	443 (HTTPS)	443	443
	SIP-internal	5061 (SIP)	5061	5061
	SIP-auth	5062 (TCP)	5062	5062
	STUN-internal	3478 (UDP)	3478 (UDP)	3478 (UDP)

Edge Virtual Server IP Address 2: OCS Edge Conferencing servers

Role:Pool Members	Service	Monitor Port (Monitor Type)	Pool TCP Port	VIP TCP Port
OCS R2 Edge Servers: OCS R2 Edge Conferencing Servers	HTTPS-external	443 (HTTPS)	443	443
	SIP-external	5061 (SIP)	5061	5061
	STUN-external	3478 (UDP)	3478 (TCP)	3478 (TCP)

Configuring the BIG-IP LTM for the Edge server additional protocols and services

Use the following procedures as a template for configuring the Edge server protocols and services applicable to your configuration.

Creating the Edge server health monitors

The first step is to configure the health monitor. This procedure uses entries from the **Monitor Port (Monitor Type)** column in the tables above.

To configure the OCS Edge server health monitors

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**. The Monitors screen opens.

2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a unique name for this Monitor.
We recommend using *Edge* and the *protocol* in the name to avoid confusion. For example, **ocs-edge-sip**.
4. From the **Type** list, select the monitor type found in the **Monitor Port (Monitor Type)** column in parenthesis. For example, if the column contains **5061 (SIP)**, select **SIP** from the **Type** list.
The Monitor configuration options appear.
5. From the **Configuration** list, select **Advanced**.
The advanced configuration options appear.
6. In the Configuration section, in the **Interval** and **Timeout** boxes, type an Interval and Timeout. We recommend at least a 1:3 +1 ratio between the interval and the timeout. In our example, we use a Interval of **30** and a Timeout of **91**.
7. In the **Alias Service Port** box, type the appropriate port found in the **Monitor Port (Monitor Type)** column. For example, if the column contains **5061 (SIP)**, type **5061** in the **Alias Service Port** box.
8. All other configuration settings are optional, configure as applicable for your deployment.
9. Click the **Finished** button.

Creating the Edge server pools

The next step is to create the pool on the BIG-IP LTM system. This procedure uses entries from the **Pool TCP Port** column in the tables above.

To create the Edge server pools

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
The Pool screen opens.
2. Click the **Create** button. The New Pool screen opens.
3. From the **Configuration** list, select **Advanced**.
4. In the **Name** box, type a unique name for this Pool.
We recommend using *Edge* and the *protocol* in the name to avoid confusion. For example, **ocs-edge-sip**.
5. In the Health Monitors section, from the **Available** list, select the name of the monitor you created in *Creating the Edge server health monitors*, on page 7, and click the Add (<<) button. Be sure to use the monitor that is associated with the same service or protocol.
6. *Optional:* In the **Slow Ramp Time** box, type **300**.
Because we are using the Least Connections load balancing method, we set the Slow Ramp Time in order to ensure that if a pool member becomes available after maintenance or a new member is added, the BIG-IP does not send all new connections to that member (a newly available member always has the least number of connections).

-
7. In the Resources section, from the **Load Balancing Method** list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network). We recommend selecting **Least Connections (node)** for all pools.
 8. In the New Members section, you add the Office Communications Servers to the pool.
 - a) In the **Address** box, type the IP address of the Office Communications Edge Server.
 - b) In the **Service Port** box, type the service number from the **Pool TCP Port** column in the table above. For example, if you are configuring the SIP-auth pool, use port **5062**.
 - c) Click the **Add** button to add the member to the list.
 - d) Repeat steps a-c for each Office Communications Edge Server you want to add to the pool.
 9. Click the **Finished** button.

Creating the Edge server TCP profiles

The next step is to create a TCP profile. For the Edge servers, most of your end users are likely connecting over the wide area network (WAN), so we recommend using the **tcp-wan-optimized** parent profile.

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, from the **Protocol** menu, select **TCP**.
3. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a unique name for this profile. We recommend using *Edge*, the *protocol*, and the type of TCP profile in the name to avoid confusion. For example, **ocs-edge-sip-wan-opt**.
5. From the **Parent Profile** list, select **tcp-wan-optimized**.
6. In the **Idle Timeout** row, check the **Custom** box. Leave the list set to **Specify**, and in the box, type **1800**.
7. Modify the rest of the settings as applicable for your network. The default settings should suffice for most networks.
8. Click the **Finished** button.

Creating the Edge server virtual servers

The final step in this section is to define a virtual server that references the profile and pool you created. This procedure uses entries from the **VIP TCP Port** column in the tables above.

To create the Edge virtual servers

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. Click the **Create** button. The New Virtual Server screen opens.
3. In the **Name** box, type a unique name for this Virtual Server. We recommend using *Edge* and the *protocol* in the name to avoid confusion. For example, **ocs-edge-sipauth-vs**.
4. In the Destination section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we use **10.133.36.51**.
6. In the **Service Port** box, type the service number from the **VIP TCP Port** column in the table above. For example, if you are configuring the SIP-auth virtual server, use port **5062**.
7. From the **Configuration** list, select **Advanced**.
8. From the **Protocol Profile (Client)** list, select the name of the profile you created in *Creating the Edge server TCP profiles*, on page 3-9.
9. In the Resources section, from the **Default Pool** list, select the pool you created in *Creating the Edge server pools*, on page 3-8.
10. Configure any other settings as appropriate for your configuration.
11. Click the **Finished** button.

Repeat this section for all of your Edge server protocols and services.

If you are using a BIG-IP LTM redundant system, see *Synchronizing the BIG-IP configuration if using a redundant system*, on page 1-23.

This completes the BIG-IP LTM and Microsoft Office Communications Server 2007 R2 deployment guide. To provide feedback on this deployment guide or other F5 solution documents, contact us at solutionsfeedback@f5.com.