

IMPORTANT: This guide has been archived. While the content in this guide is still valid for the products and version listed in the document, it is no longer being updated and may refer to F5 or 3rd party products or versions that have reached end-of-life or end-of-support. See <https://support.f5.com/csp/article/K11163> for more information.



Deploying F5 for Microsoft Office Web Apps Server 2013

Welcome to the F5 - Microsoft® Office Web Apps Server deployment guide. This document contains guidance on configuring the BIG-IP® Local Traffic Manager™ (LTM) and Application Acceleration Manager (AAM) for high availability and optimization of Microsoft Office Web Apps Server.

Office Web Apps is the online companion to Office Word, Excel, PowerPoint, and OneNote applications. It enables users, regardless of location, to view and edit documents. Office Web Apps gives users a browser-based viewing and editing experience by providing a representation of an Office document in the browser.

For more information on Microsoft Office Web Apps server, see <http://technet.microsoft.com/en-us/library/ff431685> or <http://office.microsoft.com/en-us/web-apps/>

This document is meant for organizations who have existing F5 deployments (or are in the process of deploying F5) for Microsoft Exchange Server 2013, Microsoft SharePoint 2013, or Microsoft Lync Server 2013, and want to use the BIG-IP system for the associated Office Web Apps implementation.

For more information on the BIG-IP system, see <http://www.f5.com/products/big-ip/>.

For F5 deployment guides on the other Microsoft applications mentioned in this document, see: www.f5.com/products/documentation/deployment-guides.view.solutions.base-application.microsoft.html

Visit the Microsoft page of F5's online developer community, DevCentral, for Microsoft forums, solutions, blogs and more: <http://devcentral.f5.com/Microsoft/>.

Products and applicable versions

Product	Versions
BIG-IP LTM, AAM	v11.4 - 12.1
Microsoft Office Web Apps	2013
iApp template version	0.1.0
Deployment Guide version	2.7 (see Document Revision History on page 36)
Last updated	05-18-2016

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

If you are looking for older versions of this or other deployment guides, check the Deployment Guide Archive tab at: <https://f5.com/solutions/deployment-guides/archive-608>

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

Contents

What is F5 iApp?	3
Prerequisites and configuration notes	3
Optional modules	4
Configuration scenarios	5
Using this guide	7
Preparing to use the iApp	8
Configuring the BIG-IP iApp for Microsoft Office Web Apps	9
Downloading and importing the Office Web Apps iApp from DevCentral	9
Advanced options	9
Template Options	10
Network	10
SSL Encryption	13
Virtual Server and Pools	14
Delivery Optimization	17
Server offload	19
Application Health	20
iRules	21
Statistics and Logging	22
Finished	22
Next steps	23
Modifying DNS settings to use the BIG-IP virtual server address	23
Appendix: Manual configuration table	24
Adding Office Web Apps support to a SharePoint 2013 virtual server	26
Creating the health monitor and pool for the Office Web Apps servers	26
Creating the iRule	26
Adding the iRule to the SharePoint 2013 virtual server	27
BIG-IP Access Policy considerations for Office Web Apps server	27
Adding multiple host domains to the Access Profile	27
Creating the internal virtual server on the BIG-IP system	28
Troubleshooting	30
Glossary	32
Document Revision History	35

What is F5 iApp?

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 Office Web Apps acts as the single-point interface for building, managing, and monitoring these servers.

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

Prerequisites and configuration notes

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

- This configuration assumes that you have correctly followed the Office Web Apps configuration steps described in <http://technet.microsoft.com/en-us/library/jj219455.aspx>
- To support the termination of SSL connections at BIG-IP system (SSL offload), you must configure the Office Web Apps farm with the **-AllowHTTP** and **-SSLOffloaded** options set to **True**. For specific instructions, see the Microsoft documentation.
- Additionally, you must have correctly configured the Microsoft application that is using Office Web Apps. The instructions for each application are located here:
 - » Microsoft Exchange Server 2013:
<http://technet.microsoft.com/en-us/library/jj150495.aspx>
 - » Microsoft SharePoint Server 2013:
<http://technet.microsoft.com/en-us/library/ff431687.aspx>
 - » Microsoft Lync Server 2013:
<http://technet.microsoft.com/library/3370ab55-9949-4f32-b88b-5cffed6aaad8>

After confirming that Office Web Apps Server is properly configured and that you can access the discovery URL from the Office Web Apps server(s), you can continue with the BIG-IP configuration.

- There are three configuration options described in this guide:
 - » **Creating a separate virtual server for Office Web Apps**
Creating a separate BIG-IP virtual server for Office Web Apps is recommended for Microsoft Exchange Server 2013 and Lync Server 2013 deployments, and can also be used for Microsoft SharePoint 2013. You can use the iApp template for this option, or configure the BIG-IP system manually.
 - » **Using an existing SharePoint 2013 virtual server for Office Web Apps**
This option is only available if you are configuring the BIG-IP system for SharePoint 2013 and Office Web Apps. This option requires creating an iRule to forward Office Web Apps traffic to the correct pool of servers, and adding the iRule to the existing SharePoint 2013 virtual server on the BIG-IP system.
See *Adding Office Web Apps support to a SharePoint 2013 virtual server on page 27*.
 - » **Modifying the BIG-IP configuration if using Access Policy Manager**
If you are using the BIG-IP Access Policy Manager (APM), there are additional modifications you must make to the BIG-IP configuration. See *BIG-IP Access Policy considerations for Office Web Apps server on page 28*.
- If you are deploying Office Web Apps to the same virtual server that receives application traffic, the SSL certificate must contain the Office Web Apps farm host name and individual server FQDNs in the Subject Alternative Name field, or it must be a wildcard certificate.
- If your SharePoint 2013 deployment is using BIG-IP AAM, you must add the Office Web Apps host name to the Acceleration policy in the Requested Hosts field. How you add the host name depends on how you configured the BIG-IP system:
 - » *If you used the BIG-IP iApp template to configure BIG-IP AAM for SharePoint:*
From the Application Service Properties page, on the Menu bar, click **Reconfigure**. In the Protocol Optimization section, find the question that asks for the FQDNs end users use to access SharePoint. Click **Add** and then type the FQDN for the Office Web Apps farm. Click **Finished**.

- » *If you configured BIG-IP AAM for SharePoint manually:*
On the Main tab, expand **Acceleration** and then click **Web Application**. Click the SharePoint Application, and then click **Add Host**. Type the host name for the Office Web Apps farm and then click **Save**.
- If you are using the BIG-IP AAM for Symmetric optimization between two BIG-IP systems (optional), you must have pre-configured the BIG-IP AAM for Symmetric Optimization using the Quick Start wizard or manually configured the necessary objects. See the BIG-IP AAM documentation (<http://support.f5.com/kb/en-us/products/big-ip-aam.html>) for specific instructions on configuring BIG-IP AAM for Symmetric Optimization.

Skip ahead Advanced

If you are already familiar with the iApp or the BIG-IP system, you can skip the Configuration Scenario and Preparation sections. See:

- [Configuring the BIG-IP iApp for Microsoft Office Web Apps on page 9](#) if using the iApp template, or
- [Appendix: Manual configuration table on page 25](#) if configuring the BIG-IP system manually.

Optional modules

This iApp allows you to use two optional modules on the BIG-IP system: Application Visibility Reporting (AVR) and Application Acceleration Manager (AAM). To take advantage of these modules, they must be licensed and provisioned before starting the iApp template. For more information on licensing modules, contact your sales representative.

- **BIG-IP AAM** (formerly BIG-IP WAN Optimization Manager and WebAccelerator)
BIG-IP AAM provides application, network, and front-end optimizations to ensure consistently fast performance for today's dynamic web applications, mobile devices, and wide area networks. With sophisticated execution of caching, compression, and image optimization, BIG-IP AAM decreases page download times. You also have the option of using BIG-IP AAM for symmetric optimization between two BIG-IP systems. For more information on BIG-IP Application Acceleration Manager, see <http://www.f5.com/products/big-ip/big-ip-application-acceleration-manager/overview/>.
- **Application Visibility and Reporting**
F5 Analytics (also known as Application Visibility and Reporting or AVR) is a module on the BIG-IP system that lets customers view and analyze metrics gathered about the network and servers as well as the applications themselves. Making this information available from a dashboard-type display, F5 Analytics provides customized diagnostics and reports that can be used to optimize application performance and to avert potential issues. The tool provides tailored feedback and recommendations for resolving problems. Note that AVR is licensed on all systems, but must be provisioned before beginning the iApp template.

Configuration scenarios

With the iApp template for Office Web Apps, you can configure the BIG-IP system to optimize and direct traffic to the servers with ease. You can also configure the BIG-IP system for different system scenarios using the options found in the iApp, as described in this section.

Configuring the BIG-IP system as reverse (or inbound) proxy

In its traditional role, the BIG-IP system is a reverse proxy. The system is placed in the network between the clients and the servers. Incoming requests are handled by the BIG-IP system, which interacts on behalf of the client with the desired server or service on the server. This allows the BIG-IP system to provide scalability, availability, server offload, and much more, all completely transparent to the client.




Figure 1: Using the BIG-IP system as a reverse proxy

To configure this scenario

There are no questions in the iApp template that you must answer in a specific way for the BIG-IP system to act as a reverse proxy, the BIG-IP system acts as a reverse proxy by default.

Accelerating application traffic over the WAN

The iApp enables you to use the BIG-IP system's Application Acceleration Manager module to optimize and secure your web traffic over the WAN (wide area network). The iApp uses the default *iSession profile* to create a secure tunnel between BIG-IP systems to accelerate and optimize the traffic.

In this scenario, you must have a symmetric BIG-IP deployment (as shown in Figure 2), with a BIG-IP system between your clients and the WAN, and another between the WAN and your servers. You run the iApp template on each of the BIG-IP systems, using the settings found in the following table.




Figure 2: Using an iSession tunnel to secure and optimize traffic between two BIG-IP systems

To configure this scenario

If you select this option, you must have already configured the BIG-IP AAM for Symmetric Optimization as mentioned in the prerequisites. See the BIG-IP AAM documentation available on AskF5™ (<http://support.f5.com/kb/en-us/products/big-ip-aam.html>) for specific instructions on configuring BIG-IP AAM for Symmetric Optimization.

To configure the system for this scenario, at a minimum you must answer the following questions with the appropriate answers in the iApp template as shown in the following table.

The table assumes you are configuring the BIG-IP system on the client side of the WAN.

iApp template question	Your answer
On the BIG-IP system between <i>clients</i> and the WAN	
What type of network connects clients to the BIG-IP system? (on page 10)	LAN or WAN as appropriate
What type of network connects servers to the BIG-IP system? (on page 11)	WAN through another BIG-IP system
Do you want to create a new pool or use an existing one?	Typically you would leave this at the default for this scenario (Do not use a pool), however you could create a pool of local servers to use as a fallback in case the WAN becomes unavailable.
On the BIG-IP system between <i>servers</i> and the WAN	
What type of network connects clients to the BIG-IP system? (on page 10)	WAN through another BIG-IP system
What type of network connects servers to the BIG-IP system? (on page 11)	LAN or WAN as appropriate (Typically LAN)

Using the BIG-IP system with SSL traffic

The Office Web Apps iApp template provides three different options for dealing with encrypted traffic: SSL Offload, SSL Bridging, and encrypting previously unencrypted traffic to the servers. There is also an option if you do not need the system to process SSL traffic.

- SSL Offload**
 When performing SSL offload, the BIG-IP system accepts incoming encrypted traffic, decrypts (or terminates) it, and then sends the traffic to the servers unencrypted. By saving the servers from having to perform the decryption duties, F5 improves server efficiency and frees server resources for other tasks. Certificates and keys are stored on the BIG-IP system.
- SSL Bridging**
 With SSL Bridging, also known as SSL re-encryption, the BIG-IP system accepts incoming encrypted traffic, decrypts it for processing, and then re-encrypts the traffic before sending it back to the servers. This is useful for organizations that have requirements for the entire transaction to be SSL encrypted. In this case, SSL certificates and keys must be stored and maintained on the BIG-IP system and the servers.
- SSL pass-through**
 With SSL pass-through, the BIG-IP system does not process the encrypted traffic at all, just sends it on to the servers.
- No SSL (plaintext)**
 In this scenario, the BIG-IP system does not perform any SSL processing, as all traffic is only plaintext.
- Server-side encryption**
 In this scenario, the BIG-IP system accepts unencrypted traffic and then encrypts it before sending it to the servers. While more uncommon than offload or bridging, it can be useful for organizations requiring all traffic behind the system to be encrypted.




Figure 3: SSL options

To configure these scenarios

For SSL offload or SSL bridging, you must have imported a valid SSL certificate and key onto the BIG-IP system. Importing certificates and keys is not a part of the template, see **System > File Management > SSL Certificate List**, and then click **Import**.

iApp template question	Your answer
How should the BIG-IP system handle SSL traffic (on page 13)	Select the appropriate option for your configuration: SSL Offload: Encrypt to clients, plaintext to servers SSL Bridging: Terminate SSL from clients, re-encrypt to servers SSL Pass-Through: Encrypted traffic is forwarded without decryption No SSL: Plaintext to clients and servers Server-side encryption: Plaintext to clients, encrypt to servers

Using this guide

This deployment guide is intended to help users deploy web-based applications 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 guide to configure the iApp template

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

The iApp template configuration portion of this guide walks you through the entire iApp, giving detailed information not found in the iApp or inline 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
 - **sub choice (and so on)**

Advanced options/questions in the template are marked with the Advanced icon: **Advanced**. These questions only appear if you select the Advanced configuration mode.

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 Office Web Apps 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 on page 25*.

Preparing to use the iApp

In order to use the iApp, it is helpful to have some information, such as server IP addresses and domain information before you begin. Use the following table for information you may need to complete the template. The table does not contain every question in the template, but rather includes the information that is helpful to have in advance. More information on specific template questions can be found on the individual pages.

BIG-IP System Preparation Table																			
Basic/Advanced mode	In the iApp, you can configure the system for your application with F5 recommended settings (Basic mode) which are a result of extensive testing and tuning with a wide variety of applications. Advanced mode allows configuring the BIG-IP system on a much more granular level, configuring specific options, or using your own pre-built profiles or iRules. Basic/Advanced "configuration mode" is independent from the Basic/Advanced list at the very top of the template which only toggles the Device and Traffic Group options (see page 9)																		
Network	<table border="1"> <thead> <tr> <th colspan="3">Type of network between clients and the BIG-IP system</th> <th colspan="3">Type of network between servers and the BIG-IP system</th> </tr> </thead> <tbody> <tr> <td>LAN</td> <td>WAN</td> <td>WAN through another BIG-IP system</td> <td>LAN</td> <td>WAN</td> <td>WAN through another BIG-IP system</td> </tr> </tbody> </table> <p>If WAN through another BIG-IP system, you must have BIG-IP AAM pre-configured for Symmetric Optimization.</p>	Type of network between clients and the BIG-IP system			Type of network between servers and the BIG-IP system			LAN	WAN	WAN through another BIG-IP system	LAN	WAN	WAN through another BIG-IP system						
	Type of network between clients and the BIG-IP system			Type of network between servers and the BIG-IP system															
	LAN	WAN	WAN through another BIG-IP system	LAN	WAN	WAN through another BIG-IP system													
	<table border="1"> <thead> <tr> <th colspan="2">Where are BIG-IP virtual servers in relation to the servers</th> <th colspan="2">Expected number of concurrent connections per server</th> </tr> </thead> <tbody> <tr> <td>Same subnet</td> <td>Different subnet</td> <td>More than 64k concurrent</td> <td>Fewer than 64k concurrent</td> </tr> </tbody> </table> <p>If they are on different subnets, you need to know if the servers have a route through the BIG-IP system. If there is not a route, you need to know the number of concurrent connections.</p> <p>If more than 64k per server, you need an available IP address for each 64k connections you expect for the SNAT Pool</p>	Where are BIG-IP virtual servers in relation to the servers		Expected number of concurrent connections per server		Same subnet	Different subnet	More than 64k concurrent	Fewer than 64k concurrent										
Where are BIG-IP virtual servers in relation to the servers		Expected number of concurrent connections per server																	
Same subnet	Different subnet	More than 64k concurrent	Fewer than 64k concurrent																
SSL Encryption	<table border="1"> <thead> <tr> <th>SSL offload or SSL bridging</th> <th>Re-encryption (Bridging and server-side encryption)</th> </tr> </thead> <tbody> <tr> <td> <p>If configuring the system for SSL Offload or SSL Bridging, you must have imported a valid SSL certificate and key onto the BIG-IP system. You have the option of also using an Intermediate (chain) certificate as well if required in your implementation.</p> <p>Certificate:</p> <p>Key:</p> <p>Intermediate Certificate (optional):</p> </td> <td> <p>When the BIG-IP system encrypts traffic to the servers, it is acting as an SSL client and by default we assume the servers do not expect the system to present its client certificate on behalf of clients traversing the virtual server. If your servers expect the BIG-IP system to present a client certificate, you must create a custom Server SSL profile outside of the template with the appropriate certificate and key.</p> </td> </tr> </tbody> </table>	SSL offload or SSL bridging	Re-encryption (Bridging and server-side encryption)	<p>If configuring the system for SSL Offload or SSL Bridging, you must have imported a valid SSL certificate and key onto the BIG-IP system. You have the option of also using an Intermediate (chain) certificate as well if required in your implementation.</p> <p>Certificate:</p> <p>Key:</p> <p>Intermediate Certificate (optional):</p>	<p>When the BIG-IP system encrypts traffic to the servers, it is acting as an SSL client and by default we assume the servers do not expect the system to present its client certificate on behalf of clients traversing the virtual server. If your servers expect the BIG-IP system to present a client certificate, you must create a custom Server SSL profile outside of the template with the appropriate certificate and key.</p>														
SSL offload or SSL bridging	Re-encryption (Bridging and server-side encryption)																		
<p>If configuring the system for SSL Offload or SSL Bridging, you must have imported a valid SSL certificate and key onto the BIG-IP system. You have the option of also using an Intermediate (chain) certificate as well if required in your implementation.</p> <p>Certificate:</p> <p>Key:</p> <p>Intermediate Certificate (optional):</p>	<p>When the BIG-IP system encrypts traffic to the servers, it is acting as an SSL client and by default we assume the servers do not expect the system to present its client certificate on behalf of clients traversing the virtual server. If your servers expect the BIG-IP system to present a client certificate, you must create a custom Server SSL profile outside of the template with the appropriate certificate and key.</p>																		
Virtual Server and Pools	<table border="1"> <thead> <tr> <th>Virtual server</th> <th>Office Web Apps server pool</th> </tr> </thead> <tbody> <tr> <td> <p>The <i>virtual server</i> is the address clients use to access the servers.</p> <p>IP address for the virtual server:</p> <p>Associated service port:</p> <p>FQDN clients will use to access the servers:</p> </td> <td> <p>The <i>load balancing pool</i> is the LTM object that contains the servers.</p> <p>IP addresses of the servers:</p> <table border="1"> <tbody> <tr><td>1:</td><td>3:</td></tr> <tr><td>2:</td><td>4:</td></tr> <tr><td>3:</td><td>5:</td></tr> <tr><td>4:</td><td>6:</td></tr> <tr><td>5:</td><td>7:</td></tr> <tr><td>6:</td><td>8:</td></tr> <tr><td>7:</td><td>9:</td></tr> </tbody> </table> </td> </tr> </tbody> </table>	Virtual server	Office Web Apps server pool	<p>The <i>virtual server</i> is the address clients use to access the servers.</p> <p>IP address for the virtual server:</p> <p>Associated service port:</p> <p>FQDN clients will use to access the servers:</p>	<p>The <i>load balancing pool</i> is the LTM object that contains the servers.</p> <p>IP addresses of the servers:</p> <table border="1"> <tbody> <tr><td>1:</td><td>3:</td></tr> <tr><td>2:</td><td>4:</td></tr> <tr><td>3:</td><td>5:</td></tr> <tr><td>4:</td><td>6:</td></tr> <tr><td>5:</td><td>7:</td></tr> <tr><td>6:</td><td>8:</td></tr> <tr><td>7:</td><td>9:</td></tr> </tbody> </table>	1:	3:	2:	4:	3:	5:	4:	6:	5:	7:	6:	8:	7:	9:
Virtual server	Office Web Apps server pool																		
<p>The <i>virtual server</i> is the address clients use to access the servers.</p> <p>IP address for the virtual server:</p> <p>Associated service port:</p> <p>FQDN clients will use to access the servers:</p>	<p>The <i>load balancing pool</i> is the LTM object that contains the servers.</p> <p>IP addresses of the servers:</p> <table border="1"> <tbody> <tr><td>1:</td><td>3:</td></tr> <tr><td>2:</td><td>4:</td></tr> <tr><td>3:</td><td>5:</td></tr> <tr><td>4:</td><td>6:</td></tr> <tr><td>5:</td><td>7:</td></tr> <tr><td>6:</td><td>8:</td></tr> <tr><td>7:</td><td>9:</td></tr> </tbody> </table>	1:	3:	2:	4:	3:	5:	4:	6:	5:	7:	6:	8:	7:	9:				
1:	3:																		
2:	4:																		
3:	5:																		
4:	6:																		
5:	7:																		
6:	8:																		
7:	9:																		
Profiles	<p>The iApp template can create <i>profiles</i> using the F5 recommended settings, or you can choose Do not use many of these profiles). F5 recommends using the profiles created by the iApp; however you also have the option of creating your own custom profile outside the iApp and selecting it from the list. The iApp gives the option of selecting the following profiles (some only in Advanced mode). Any profiles must be present on the system before you can select them in the iApp.</p> <p>HTTP Persistence HTTP Compression TCP LAN TCP WAN OneConnect Web Acceleration NTLM iSession</p>																		
Health monitor	<table border="1"> <thead> <tr> <th>HTTP request</th> <th>User account</th> </tr> </thead> <tbody> <tr> <td> <p>In Advanced mode, you have the option of selecting the type of HTTP request the health monitor uses: GET or POST.</p> <p>You can also specify Send and Receive strings to more accurately determine server health.</p> <p>Send string (the URI sent to the servers):</p> <p>Receive string (what the system expects in return):</p> <p>POST Body (only if using POST):</p> </td> <td> <p>Also in advanced mode, the monitor can attempt to authenticate to the servers as a part of the health check. If you want the monitor to require credentials, create a user account specifically for this monitor that has no additional permissions and is set to never expire. Account maintenance becomes a part of the health monitor, as if the account is deleted or otherwise changed, the monitor will fail and the servers will be marked down.</p> </td> </tr> </tbody> </table>	HTTP request	User account	<p>In Advanced mode, you have the option of selecting the type of HTTP request the health monitor uses: GET or POST.</p> <p>You can also specify Send and Receive strings to more accurately determine server health.</p> <p>Send string (the URI sent to the servers):</p> <p>Receive string (what the system expects in return):</p> <p>POST Body (only if using POST):</p>	<p>Also in advanced mode, the monitor can attempt to authenticate to the servers as a part of the health check. If you want the monitor to require credentials, create a user account specifically for this monitor that has no additional permissions and is set to never expire. Account maintenance becomes a part of the health monitor, as if the account is deleted or otherwise changed, the monitor will fail and the servers will be marked down.</p>														
HTTP request	User account																		
<p>In Advanced mode, you have the option of selecting the type of HTTP request the health monitor uses: GET or POST.</p> <p>You can also specify Send and Receive strings to more accurately determine server health.</p> <p>Send string (the URI sent to the servers):</p> <p>Receive string (what the system expects in return):</p> <p>POST Body (only if using POST):</p>	<p>Also in advanced mode, the monitor can attempt to authenticate to the servers as a part of the health check. If you want the monitor to require credentials, create a user account specifically for this monitor that has no additional permissions and is set to never expire. Account maintenance becomes a part of the health monitor, as if the account is deleted or otherwise changed, the monitor will fail and the servers will be marked down.</p>																		
BIG-IP Application Acceleration Manager	You can optionally use the BIG-IP Application Acceleration Manager (AAM) module to help accelerate your HTTP traffic. To use BIG-IP AAM, it must be fully licensed and provisioned on your BIG-IP system. Consult your F5 sales representative for details. If you are using BIG-IP AAM, and want to use a custom Web Acceleration policy, it must have an Acceleration policy attached.																		
iRules	In Advanced mode, you have the option of attaching iRules you create to the virtual server created by the iApp. For more information on iRules, see https://devcentral.f5.com/irules . Any iRules you want to attach must be present on the system at the time you are running the iApp.																		

Configuring the BIG-IP iApp for Microsoft Office Web Apps

Use the following guidance to help you configure Microsoft Lync Server using the BIG-IP iApp template. You must have downloaded and imported the iApp from DevCentral before beginning.

Downloading and importing the Office Web Apps iApp from DevCentral

The first task is to download the latest iApp for Microsoft Office Web Apps from DevCentral and import it onto the BIG-IP system.

To download and import the iApp from DevCentral

1. Open a web browser and go to <https://devcentral.f5.com/wiki/iApp.Microsoft-Office-Web-Apps-iApp-v0-1-0.ashx>
2. Download the **f5.microsoft_office_web_apps_2013.<latest-version>.zip** file to a location accessible from your BIG-IP system.

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

3. Extract (unzip) the **f5.microsoft_office_web_apps_2013.<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 iApp template

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 **Office_web_apps-iapp_**.
5. From the **Template** list, select **f5.microsoft_office_web_apps_2013.<latest-version>**. The Office Web Apps template opens.

Advanced options

If you select **Advanced** from the **Template Selection** list at the top of the page, 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 product 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.

Template Options

This section contains general questions about the way you configure the iApp template.

1. **Do you want to see inline help?**

Choose 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 the 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 text**

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.

2. **Which configuration mode do you want to use?**

Select whether you want to use F5 recommended settings, or have more granular, advanced options presented.

- **Basic - Use F5's recommended settings**

In Basic configuration mode, options like load balancing method and parent profiles are all set automatically. The F5 recommended settings come as a result of extensive testing with web applications, so if you are unsure, choose Basic.

- **Advanced - Configure advanced options**

In Advanced configuration mode, you have more control over individual settings and objects, such as server-side optimizations and advanced options like Slow Ramp Time and Priority Group Activation. You can also choose to attach iRules you have previously created to the Application Service. The Advanced option provides more flexibility for experienced users.

As mentioned, advanced options in the template are marked with the Advanced icon: **Advanced**. If you are using Basic/F5 recommended settings, you can skip the questions with the Advanced icon.

Network

This section contains questions about your networking configuration.

1. **What type of network connects clients to the BIG-IP system?**

Choose the type of network that connects your clients to the BIG-IP system. If you choose WAN or LAN, the BIG-IP system uses this information to determine the appropriate TCP optimizations. If you choose WAN through another BIG-IP system, the system uses a secure optimized tunnel (called an iSession tunnel) for traffic between BIG-IP systems on opposite sides of the WAN.

Only choose this option if you have another BIG-IP system across the WAN that will be a part of this implementation.

- **Local area network (LAN)**

Select this option if most clients are connecting to the BIG-IP system on a LAN. This field is used to determine the appropriate TCP profile which is optimized for LAN clients. In this case, the iApp creates a TCP profile using the *tcp-lan-optimized* parent with no additional modifications.

- **Wide area network**

Select this option if most clients are connecting to the BIG-IP system over a WAN. This field is used to determine the appropriate TCP profile which is optimized for WAN clients. In this case, the iApp creates a TCP profile using the *tcp-wan-optimized* parent with no additional modifications.

- **WAN through another BIG-IP system**

Select this option if client traffic is coming to this BIG-IP system from a remote BIG-IP system across a WAN. As mentioned in the introduction to this question, the iApp creates an iSession tunnel between this BIG-IP system and the BIG-IP system you will configure (or already have configured) on the other side of the WAN.

If you select this option, you must have already initially configured the BIG-IP AAM for Symmetric Optimization. See the BIG-IP AAM documentation available on Ask F5 for specific instructions on configuring BIG-IP AAM for Symmetric Optimization.

2. **Do you want to restrict client traffic to specific VLANs?** **Advanced**

The BIG-IP system allows you to restrict client traffic to specific VLANs that are present on the system. This can provide an additional layer of security, as you can allow or deny traffic from the VLANs you choose. By default, all VLANs configured on the system are enabled. If you select to enable or disable traffic on specific VLANs, you must specify the VLANs in the next question.

The VLAN objects must already be configured on this BIG-IP system before you can select them.

- **Enable traffic on all VLANs and Tunnels**


Choose this option to allow traffic from all VLANs and Tunnels. If you select this option, the question asking about VLANs disappears. Continue with #3.

- **Yes, enable traffic only on the VLANs I specify**

Choose this option to restrict client traffic to specific VLANs that you choose in the following question. The system will accept client traffic from these VLANs, and deny traffic from all other VLANs on the system.

- a. On which VLANs should traffic be enabled or disabled?

Use this section to specify the VLANs that will accept client traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so click the VLANs and then use the Move buttons to adjust list membership.


 **Note:** *If you choose to allow traffic from certain VLANs, when additional VLANs are added to the BIG-IP system at a later time, this iApp configuration will deny traffic from these VLANs by default. To accept traffic from these VLANs, you must re-enter the template and add the VLAN(s).*

- **Yes, disable traffic only on the VLANs I specify**

Choose this option to deny client traffic from the specific VLANs that you choose in the following question. The system will refuse client traffic from these VLANs, and accept traffic from all other VLANs on the system.

- a. On which VLANs should traffic be enabled or disabled?

Use this section to specify the VLANs that should not accept client traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so it is critical in this case that you click the VLANs and then use the Move button (>>) to adjust list membership.

 **Warning** *If you choose to disable certain VLANs, you must move at least one VLAN to the Options list. Otherwise, the system will deny traffic from all VLANs on the box, and the configuration, although valid, will not pass any traffic.*

3. What type of network connects servers to the BIG-IP system?

Choose the type of network that connects your servers to the BIG-IP system. Similar to the question about clients connecting to the BIG-IP system, if you choose WAN or LAN, the system uses this information to determine the appropriate TCP optimizations. If you choose WAN through another BIG-IP system, the system uses a secure and optimized tunnel (called an iSession tunnel) for traffic between BIG-IP systems on opposite sides of the WAN. Only choose this option if you have another BIG-IP system across the WAN that will be a part of this implementation.

- **Local area network (LAN)**

Select this option if the servers connect to the BIG-IP system on a LAN. This field is used to determine the appropriate TCP profile. In this case, the iApp creates a TCP profile using the *tcp-lan-optimized* parent with no additional modifications.

- **Wide area network**

Select this option if the servers connect to the BIG-IP system over a WAN. This field is used to determine the appropriate TCP profile. In this case, the iApp creates a TCP profile using the *tcp-wan-optimized* parent with no additional modifications.

- **WAN through another BIG-IP system**

Select this option if servers are across a WAN behind another BIG-IP system. As mentioned in the introduction to this question, the iApp creates an iSession tunnel between this BIG-IP system and the BIG-IP system you will configure (or already have configured) on the other side of the WAN.

If you select this option, you must have already initially configured the BIG-IP AAM for Symmetric Optimization. See the BIG-IP AAM documentation available on Ask F5 for specific instructions on configuring BIG-IP AAM for Symmetric Optimization.

4. Where will your BIG-IP virtual servers be in relation to the Office Web Apps servers?

Select whether your BIG-IP virtual servers are on the same subnet as your servers, or on different subnets. This setting is used to determine the secure NAT ([SNAT](#)) and routing configuration.

- **BIG-IP virtual server IP and Office Web Apps servers are on the same subnet**

If the BIG-IP virtual servers and Office Web Apps servers are on the same subnet, SNAT is configured on the BIG-IP virtual server and you must specify the number of concurrent connections.

a. How many connections to you expect to each Office Web Apps server?

Select whether you expect more or fewer than 64,000 concurrent connections to each server. This answer is used to determine what type of SNAT that system uses. A SNAT is an object that maps the source client IP address in a request to a translation address defined on the BIG-IP device. The system configures SNAT Auto Map (fewer than 64,000) or a SNAT pool (more than 64,000).

- **Fewer than 64,000 concurrent connections**

Select this option if you expect fewer than 64,000 concurrent connections per server. With this option, the system applies SNAT Auto Map, doesn't require any additional IP addresses, as an existing self IP address is used for translation. Continue with *Virtual Server and Pools on page 15*.

- **More than 64,000 concurrent connections**

Select this option if you have a very large deployment and expect more than 64,000 connections at one time. The iApp creates a SNAT Pool, for which you need one IP address for each 64,000 connections you expect.

a. Create a new SNAT pool or use an existing one?

If you have already created a SNAT pool on the BIG-IP system, you can select it from the list. Otherwise, the system creates a new SNAT pool with the addresses you specify.

- **Create a new SNAT pool**


Select this option (the default) to enable the system to create a new SNAT pool. You must specify the appropriate number of IP addresses in the next question.

a. Which IP addresses do you want to use for the SNAT pool?

Specify one otherwise unused IP address for every 64,000 concurrent connections you expect, or fraction thereof. Click **Add** for additional rows. Do not use any BIG-IP self IP addresses.

- **Select a SNAT pool**

Select the SNAT pool you created for this deployment from the list.

 **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.*

- **BIG-IP virtual servers and Office Web Apps servers are on different subnets**

If the BIG-IP virtual servers and servers are on different subnets, the following question appears.

a. How have you configured routing on your Office Web Apps servers?

If you chose different subnets, this question appears asking whether the servers use this BIG-IP system's self IP address as their default gateway. Select the appropriate answer.

- **Servers have a route to clients through the BIG-IP system**

Choose this option if the servers use the BIG-IP system as their default gateway. In this case, no configuration is needed to support your environment to ensure correct server response handling. Continue with the next section.

- **Servers do not have a route to clients through the BIG-IP system**

If the servers do not use the BIG-IP system as their default gateway, [SNAT](#) is configured on the BIG-IP virtual server and you must select the expected number of concurrent connections in the next question.

a. How many connections to you expect to each Office Web Apps server?

Select whether you expect more or fewer than 64,000 concurrent connections to each server. This answer is used to determine what type of SNAT that system uses. A SNAT is an object that maps the source client IP address in a request to a translation address defined on the BIG-IP device. The system configures SNAT Auto Map (fewer than 64,000) or a SNAT pool (more than 64,000).

- **Fewer than 64,000 concurrent connections**

Select this option if you expect fewer than 64,000 concurrent connections per server. With this option, the system applies SNAT Auto Map, doesn't require any additional IP addresses, as an existing self IP address is used for translation. Continue with the *SSL Encryption* section.

- **More than 64,000 concurrent connections**

Select this option if you have a very large deployment and expect more than 64,000 connections at one time. The iApp creates a SNAT Pool, for which you need one IP address for each 64,000 connections you expect.

a. Create a new SNAT pool or use an existing one?

If you have already created a SNAT pool on the BIG-IP system, you can select it from the list. Otherwise, the system creates a new SNAT pool with the addresses you specify.

- **Create a new SNAT pool**
Select this option (the default) to enable the system to create a new SNAT pool. You must specify the appropriate number of IP addresses in the next question.
 - a. *Which IP addresses do you want to use for the SNAT pool?*
Specify one otherwise unused IP address for every 64,000 concurrent connections you expect, or fraction thereof. Click **Add** for additional rows. Do not use any self IP addresses on the BIG-IP system.
 - **Select a SNAT pool**
Select the SNAT pool you created for this deployment from the list.
- i 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.*

SSL Encryption

Before running the template you should have already imported a certificate and key onto the BIG-IP system. While the BIG-IP system does include a self-signed SSL certificate that can be used internally or for testing, we strongly recommend importing a certificate and key issued from a trusted Certificate Authority for processing client-side SSL.

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

1. **How should the BIG-IP system handle SSL traffic?**

There are four options for configuring the BIG-IP system for SSL traffic. Select the appropriate mode for your configuration.

- **Encrypt to clients, plain text to servers (SSL Offload)**
Choose this method if you want the BIG-IP system to offload SSL processing from the servers. You need a valid SSL certificate and key for this method.
 - a. *Which Client SSL profile do you want to use?* **Advanced**
Select whether you want the iApp to create a new Client SSL [profile](#), or if you have already created a Client SSL profile which contains the appropriate SSL certificate and key.

Unless you have requirements for configuring specific Client SSL settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic > Profiles > SSL > Client** to create a Client SSL profile. To select any new profiles you create, you need to restart or reconfigure this template.
 - **Select an existing Client SSL profile**
If you created a Client SSL profile for this implementation, select it from the list. If you select an existing Client SSL profile, the rest of the questions in this section disappear. Continue with *Virtual Server and Pools on page 15*.
 - **Create a new Client SSL profile**
Select this option for the iApp to create a new Client SSL profile.
 - 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.
- **Terminate SSL from clients, re-encrypt to servers (SSL Bridging)**
Choose this method if you want the BIG-IP system to terminate SSL to process it, and then re-encrypt the traffic to the servers (SSL Bridging). You need a valid SSL certificate and key for the client-side, and optionally for the server-side (see #b).

a. Which Client SSL profile do you want to use? **Advanced**

Select whether you want the iApp to create a new Client SSL [profile](#), or if you have already created a Client SSL profile which contains the appropriate SSL certificate and key.

Unless you have requirements for configuring specific Client SSL settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic > Profiles > SSL > Client** to create a Client SSL profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Select an existing Client SSL profile**
If you created a Client SSL profile for this implementation select it from the list. If you select an existing Client SSL profile, the rest of the questions in this section disappear. Continue with *Virtual Server and Pools on page 15*.

- **Create a new Client SSL profile**

Select this option for the iApp to create a new Client SSL profile

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 implementation 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.

b. Which Server SSL profile do you want to use?

Select whether you want the iApp to create the F5 recommended Server SSL profile, or if you want to choose a Server SSL profile you already created. In this scenario, the BIG-IP system is acting as an SSL client and by default, we assume the servers do not expect the BIG-IP system to present its client certificate on behalf of clients traversing the virtual server. If your servers expect the BIG-IP system to present a client certificate, you must create a custom Server SSL profile with the appropriate certificate and key.

The default, F5 recommended Server SSL profile uses the `serverssl` parent profile. For information about the ciphers used in the Server SSL profile, see <http://support.f5.com/kb/en-us/solutions/public/8000/800/sol8802.html>.

- **Encrypted traffic is forwarded without decryption (SSL pass-through)**

Choose this method if you do not want the BIG-IP system to do anything with encrypted traffic and simply send it to the servers. This is similar to SSL bridging, although in this case the system does not decrypt then re-encrypt the traffic, it only sends it on to the servers without modification.

If you select this option, the system changes the default persistence option from Cookie to Source Address Persistence.

- **Plain text to clients, encrypt to servers**

Choose this method if you want the BIG-IP system to accept plain text from the clients and then encrypt it before sending it to the servers.

Unless you have requirements for configuring specific Server SSL settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic > Profiles > SSL > Server** to create a Server SSL profile. To select any new profiles you create, you need to restart or reconfigure this template.

The default, F5 recommended Server SSL profile uses the `serverssl` parent profile. For information about the ciphers used in the Server SSL profile, see <http://support.f5.com/kb/en-us/solutions/public/8000/800/sol8802.html>.

- **Plain text to both clients and servers**

Choose this method if the BIG-IP system is not sending or receiving any SSL traffic in this implementation.

Virtual Server and Pools

This section gathers information about your deployment that will be used in the BIG-IP [virtual server](#) and [load balancing pool](#).

1. ***What IP address do you want to use for the virtual server?***

Type the IP address you want to use for the BIG-IP virtual server. This is the address clients use (or a DNS entry resolves to this address) to access the Office Web Apps deployment via the BIG-IP system.

If necessary for your configuration, this can be a network address to create a network virtual server (you must specify an IP mask in the following question for a network virtual server). A network virtual server is a virtual server whose IP address has no bits set in the host portion of the IP address (that is, the host portion of its IP address is 0), allowing the BIG-IP system to direct client connections that are destined for an entire range of IP addresses, rather than for a single destination IP address. Thus, when any client connection targets a destination IP address that is in the network specified by the virtual server IP address, the system can direct that connection the pool of Office Web Apps servers.

2. ***If using a network virtual address, what is the IP mask?***

If you specified a network address for the virtual server (allowing the virtual server to handle multiple IP addresses), you must enter the full network mask representing the address range. If you specified a single address for the virtual server, leave this field blank.

3. ***What port do you want to use for the virtual server?***

Type the port number for the BIG-IP virtual server. For Office Web Apps deployments, this is typically 80 (HTTP) or 443 (HTTPS).

4. ***Which FQDNs will clients use to access the servers?***

Type each fully qualified domain name clients will use to access the Office Web Apps deployment. Click the **Add** button to insert additional rows. If you only have one FQDN, do not click Add.

5. ***Do you want to redirect inbound HTTP traffic to HTTPS?*** **Advanced**

This question only appears if you selected SSL Offload or SSL Bridging in the SSL question.

Select whether you want the BIG-IP system to automatically redirect HTTP traffic to the HTTPS virtual server. This is useful when users forget to use HTTPS when attempting to connect to the deployment.

- **Redirect HTTP to HTTPS**

Select this option to redirect HTTP traffic to HTTPS. If you select this option (the default), the BIG-IP system attaches a very small redirect iRule to the virtual server.

- a. ***From which port should traffic be redirected?***

Type the port number for the traffic that you want to redirect to HTTPS. The most common is port 80 (the default).

- **Do not redirect HTTP to HTTPS**

Select this option if you do not want to enable the automatic redirect.

6. ***Which HTTP profile do you want to use?*** **Advanced**

The HTTP [profile](#) contains settings for instructing the BIG-IP system how to handle HTTP traffic. Choose whether you want the iApp to create a new HTTP profile or if you have previously created an HTTP profile for this deployment.

Unless you have requirements for configuring specific HTTP settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic > Profiles > Services > HTTP** to create a HTTP profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Select an existing HTTP profile from the list**

If you already created an HTTP profile for this implementation, select it from the list.

- **Create a new HTTP profile (recommended)**

Select this option for the iApp to create a new HTTP profile.

- a. ***Should the BIG-IP system insert the X-Forwarded-For header?*** **Advanced**

Select whether you want the BIG-IP system to insert the X-Forwarded-For header in the HTTP header for logging purposes.

- **Insert the X-Forwarded-For header**

Select this option if you want the system to include the X-Forwarded-For header. You may have to perform additional configuration on your HTTP servers to log the value of this header. For more information on configuring logging refer to the server documentation.

- **Do not insert the X-Forwarded-For header**

Select this option if you do not want the system to include X-Forwarded-For in the HTTP header.

7. ***Which persistence profile do you want to use?*** **Advanced**

By using persistence, the BIG-IP system tracks and stores session data, such as the specific pool member that serviced a client request, ensuring client requests are directed to the same pool member throughout the life of a session or during subsequent sessions.

Unless you have requirements for configuring specific persistence settings, we recommend allowing the iApp to create a new profile. Creating a custom profile is not a part of this template; see **Local Traffic >> Profiles : Persistence** to create a persistence profile. To select any new profiles you create, you need to restart or reconfigure this template.

Select one of the following persistence options:

- **Use Cookie Persistence (recommended)** *<if you chose SSL pass-through, "(recommended)" does not appear>*
Leave this default option to have the BIG-IP system create a new cookie persistence profile (cookie insert mode). With Cookie persistence, the BIG-IP system uses an HTTP cookie stored on the client's computer to allow the client to reconnect to the same server previously visited. We recommend this method for most configurations, except for SSL pass-through.
- **Source IP Address persistence**
Select this option if you want to use the Source IP address (also known as simple) persistence. With this mode, the BIG-IP system assigns the built-in Source Address Affinity persistence type, and directs session requests to the same server based only on the source IP address. This is the recommended method if you are using SSL pass-through.
- **Do not use persistence**
If your implementation does not require persistent connections, select this option.
- **Select an existing persistence profile**
If you have previously created a persistence profile, you have the option of selecting it instead of allowing the iApp to create a new one. From the list, select an existing persistence profile. We recommend using a persistence profile that uses Cookie persistence, Insert mode.

8. ***Do you want to create a new pool or use an existing one?***

A [load balancing pool](#) is a logical set of servers, grouped together to receive and process traffic. When clients attempt to access the application via the BIG-IP virtual server, the BIG-IP system distributes requests to any of the servers that are members of that pool.

- **Select an existing pool**
If you have already created a pool for your Office Web Apps servers, you can select it from the list.
If you do select an existing pool, all of the rest of the questions in this section disappear.
- **Do not use a pool**
If you are deploying this iApp in such a way that you do not need a pool of Office Web Apps servers, select this option. If you specified the servers are connected to the BIG-IP system over the WAN through another BIG-IP system, this is the default option, as the system is sending traffic across the iSession tunnel to the other BIG-IP system to be distributed to the servers.
- **Create a new pool**
Leave this default option to create a new load balancing pool and configure specific options.

a. ***Which load balancing method do you want to use?*** **Advanced**

Specify the load balancing method you want to use for this pool. We recommend the default, **Least Connections (member)**.

b. ***Do you want to give priority to specific groups of servers?*** **Advanced**

Select whether you want to use Priority Group Activation. Priority Group Activation allows you to segment your servers into priority groups. With Priority Group Activation, the BIG-IP system load balances traffic according to the priority number you assign to the pool members. A higher number indicates higher priority. Traffic is only sent to the servers with the highest priority, unless the number of available servers in that priority group falls below the value you specify as the minimum. The BIG-IP then sends traffic to the group of servers with the next highest priority, and so on. See the BIG-IP documentation for more details.

- **Do not use Priority Group Activation (recommended)**

Select this option if you do not want to enable Priority Group Activation.

- **Use Priority Group Activation**

Select this option if you want to enable Priority Group Activation.

You must add a priority to each server in the Priority box described in #c.

- a. ***What is the minimum number of active members for each priority group?***

Specify the minimum number of servers that must be active to continue sending traffic to the priority group. If the number of active servers falls below this minimum, traffic will be sent to the group of servers with the next highest priority group number.

- c. ***Which Office Web Apps servers should be included in this pool?***

Specify the IP address(es) of your servers. If you have existing nodes on this BIG-IP system, you can select them from the list, otherwise type the addresses. You can optionally add a Connection Limit. If you enabled Priority Group Activation, you must also specify a Priority for each device. Click **Add** to include additional servers.

Delivery Optimization

In this section, you answer questions about how you want the BIG-IP system to optimize the delivery of your Office Web Apps traffic.

1. ***Use the BIG-IP Application Acceleration Manager?***

This question only appears if you have licensed and provisioned the BIG-IP Application Acceleration Manager (AAM).

Choose whether you want to use the BIG-IP Application Acceleration Manager (formerly known as WebAccelerator). BIG-IP Application Acceleration Manager helps accelerate your traffic.

- **Yes, use BIG-IP AAM (recommended)**

Select this option to enable BIG-IP AAM.

- **No, do not use BIG-IP AAM**

Select this option if you do not want to enable BIG-IP AAM at this time.

2. ***Which Web Acceleration profile do you want to use for caching?*** **Advanced**

Select whether you want the system to create a new Web Acceleration profile, or if you have already created a Web Acceleration profile for use in this deployment. The Web Acceleration profile contains the caching settings for this implementation.

Unless you have requirements for configuring specific acceleration settings (such as specific allowing/denying specific URIs), we recommend allowing the iApp to create a new profile. Creating a custom profile is not a part of this template; see **Local Traffic >> Profiles : Services : Web Acceleration** to create an acceleration profile. To select any new profiles you create, you need to restart or reconfigure this template.

Note if using BIG-IP AAM:

If you are using BIG-IP AAM, and want to select a custom Web Acceleration profile for caching you have already created, it must have an AAM application enabled, otherwise it does not appear in the list of caching profiles. If you want access to all Web Acceleration profiles on the box, then you must choose No to the use BIG-IP AAM question. Use a custom Web Acceleration profile only if you need to define specific URIs that should or should not be cached.

Note if not using BIG-IP AAM:

If you are not using BIG-IP AAM, we recommend you only use a custom Web Acceleration profile if you need to define specific URIs which should or should not be cached. You can continue with #6.

- **Create a profile based on optimized-caching (recommended)**

Leave this default option to create a new Web Acceleration profile for caching.

- **Do not use caching**

This question does not appear if you chose to enable BIG-IP AAM

Select this option if you do not want to enable caching on the BIG-IP system for this implementation.

- **Select an existing Web Acceleration profile**

If you have already created a Web Acceleration profile for your Office Web Apps servers, you can select it from the list.

3. ***Do you want to insert the X-WA-Info header?*** **Advanced**

This question only appears if you chose to enable BIG-IP AAM

The BIG-IP system can optionally insert an X-WA-Info response header that includes specific codes describing the properties and history of the object. The X-WA-Info response header is for informational and debugging purposes only and provides a way for you to assess the effectiveness of your acceleration policy rules.

By default, the AAM X-WA-info header is not included in the response from the BIG-IP system. If you choose to enable this header, you have two options, Standard and Debug. In Standard mode, the BIG-IP system inserts an HTTP header that includes numeric codes which indicate if and how each object was cached. In Debug mode, the BIG-IP system includes additional information which may help for extended troubleshooting.

- **Do not insert the header (recommended)**
Select this option if you do not want to insert the X-WA-Info header. Typically F5 recommends not inserting the header unless instructed to do so by an F5 Technical Support Engineer.
- **Insert the Standard header**
Select this option if you want to insert the Standard header. For detailed information on the numeric codes used by the header, see <http://support.f5.com/kb/en-us/solutions/public/13000/700/sol13798.html>
- **Insert the Debug header**
Select this option if you want to insert the Debug header for extended troubleshooting.

4. **Do you want to use the legacy AAM performance monitor?** **Advanced**

This question only appears if you chose to enable BIG-IP AAM

Enabling the legacy AAM performance monitor can adversely affect system performance. This monitor is primarily used for legacy AAM performance monitoring and debugging purposes, and can adversely affect system performance. The BIG-IP Dashboard provides performance graphs and statistics related to AAM.

- **Do not enable the legacy performance monitor (recommended)**
Select this option if you do not want to enable the legacy monitor.
- **Enable the legacy performance monitor**
Select this option if you want to enable the legacy performance monitor. Remember enabling this legacy monitor can impact overall system performance.
 - a. **For how many days should the BIG-IP system retain the data?**
Specify the number of days the BIG-IP system should retain the legacy performance data.

5. **Which acceleration policy do you want to use?** **Advanced**

This question only appears if you chose to enable BIG-IP AAM

Select one of the following predefined acceleration policies from the list.

- **Generic Policy - Complete**
This predefined acceleration policy is ideal for Apache HTTP servers, Microsoft Internet Information Services (IIS) web servers, WebLogic application servers, and IBM WebSphere Application Servers. HTML pages are cached and Intelligent Browser Referencing is enabled.
- **Generic Policy - Enhanced**
This predefined acceleration policy is ideal for Apache HTTP servers, Internet Information Services (IIS) web servers, WebLogic application servers, and IBM WebSphere Application Servers. HTML pages are cached and Intelligent Browser Referencing is enabled for includes.
- **Generic Policy - Extension Based.**
This predefined acceleration policy is ideal for High Performance policy for E-commerce applications that uses File Extensions instead of mime-types. This application policy is ideal if response-based matching is not required.
- **Generic Policy - Fundamental.**
This predefined acceleration policy is ideal for Apache HTTP servers, Internet Information Services (IIS) web servers, WebLogic application servers, and IBM WebSphere Application Servers. HTML pages are always proxied and Intelligent Browser Referencing is disabled.

6. **Which compression profile do you want to use?**

Compression improves performance and end user experience for Web applications that suffer from WAN latency and throughput bottlenecks. Compression reduces the amount of traffic sent to the client to complete a transaction.

Unless you have requirements for configuring specific compression settings, we recommend allowing the iApp to create a new profile. F5 recommends the default profile. Creating a custom profile is not a part of this template; see **Local Traffic > Profiles > Services > HTTP Compression** to create a compression profile. To select any new profiles you create, you need to restart or reconfigure this template.

7. ***How do you want to optimize client-side connections?*** **Advanced**

The client-side TCP profile optimizes the communication between the BIG-IP system and the client by controlling the behavior of the traffic which results in higher transfer rates, improved connection reliability and increased bandwidth efficiency.

Unless you have requirements for configuring specific TCP optimization settings, we recommend allowing the iApp to create a new profile. Creating a custom profile is not a part of this template; see **Local Traffic > Profiles > Protocol > TCP** to create a TCP profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Create the appropriate tcp-optimized profile (recommended)**
Select this option to have the system create the recommended TCP profile. The parent profile (either WAN or LAN optimized) is determined by your selection to the “What type of network connects clients to the BIG-IP system” question.
- **Select the TCP profile you created from the list**
If you created a custom TCP profile for the Office Web Apps servers, select it from the list.

Server offload

In this section, you configure the options for offloading tasks from the servers. This section only appears in Advanced mode.

1. ***Which OneConnect profile do you want to use?*** **Advanced**

OneConnect (connection pooling or multiplexing) improves server scalability by reducing load associated with concurrent connections and connection rate to Office Web Apps servers. When enabled, the BIG-IP system maintains one connection to each server which is used to send requests from multiple clients.

Unless you have requirements for configuring specific settings, we recommend allowing the iApp to create a new profile. F5 recommends the default profile. Creating a custom profile is not a part of this template; see **Local Traffic >> Profiles : Other : OneConnect** to create a OneConnect profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Create a profile based on the oneconnect parent (recommended)**
Select this option to have the system create the recommended OneConnect profile. The system uses the oneconnect parent profile with a Source Mask setting of 255.255.255.255.
- **Do not use a OneConnect profile**
Select this option if you do not require the BIG-IP system to perform connection pooling using a OneConnect profile.
- **Select the OneConnect profile you created from the list**
If you created a custom OneConnect profile for the Office Web Apps servers, select it from the list.

2. ***Which NTLM profile do you want to use?*** **Advanced**

The NTLM profile optimizes network performance when the system is processing NTLM traffic. When both an NTLM profile and a OneConnect profile are enabled, the system can take advantage of server-side connection pooling for NTLM connections.

If your environment uses NTLM, we recommend allowing the iApp to create a new profile unless you have requirements for configuring specific settings. Creating a custom profile is not a part of this template; see **Local Traffic >> Profiles : Other : NTLM** to create a NTLM profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Use F5’s recommended NTLM profile**
Select this option to have the system create the recommended NTLM profile. The system uses the ntlm parent profile.
- **Do not use NTLM (recommended)**
Select this option if you do not use NTLM authentication in your Office Web Apps implementation.
- **Select the NTLM profile you created from the list**
If you created a custom NTLM profile for the Office Web Apps servers, select it from the list.

3. ***How do you want to optimize server-side connections?*** **Advanced**

The server-side TCP profile optimizes the communication between the BIG-IP system and the server by controlling the behavior of the traffic which results in higher transfer rates, improved connection reliability and increased bandwidth efficiency.

Unless you have requirements for configuring specific TCP optimization settings, we recommend allowing the iApp to create a new profile. Creating a custom profile is not a part of this template; see **Local Traffic >> Profiles : Protocol : TCP** to create a TCP profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Create the appropriate tcp-optimized profile (recommended)**
Select this option to have the system create the recommended TCP profile. The parent profile (either WAN or LAN optimized) is determined by your selection to the “What type of network connects servers to the BIG-IP system” question.
- **Select the TCP profile you created from the list**
If you created a custom TCP profile for the Office Web Apps servers, select it from the list.

4. ***Should the BIG-IP system queue TCP requests?***

Select whether you want the BIG-IP system to queue TCP requests. TCP request queuing provides the ability to queue connection requests that exceed the capacity of connections for a pool, as determined by the connection limit. Consequently, instead of dropping connection requests that exceed the capacity of a pool, TCP request queuing enables those connection requests to reside within a queue according to defined conditions until capacity becomes available. For more information on TCP Request Queuing, see the *Preventing TCP Connection Requests From Being Dropped* chapter in the **BIG-IP Local Traffic Manager: Implementations** guide, available on AskF5.

i Important *TCP Request Queuing is an advanced feature and should be used only if you understand how it will affect your deployment, including application behavior and BIG-IP performance. If you enable TCP Request Queuing, you must have a Connection Limit set on at least one of the nodes when configuring the Address/Port for the nodes.*

- **No, do not enable TCP request queuing (recommended)**
Select this option if you do not want the BIG-IP system to queue TCP requests.
- **Yes, enable TCP request queuing**
Select this option if you want to enable TCP request queuing on the BIG-IP system.
 - What is the maximum number of TCP requests for the queue?***
Type the maximum number of requests you want to queue. We do **not** recommend using 0, which means unlimited and is only constrained by available memory.
 - How many milliseconds should requests remain in the queue?***
Type a number of milliseconds for the TCP request timeout value.

5. ***Use a Slow Ramp time for newly added servers?*** **Advanced**

With Slow Ramp, the BIG-IP system gradually adds connections to a newly-enabled or newly-added server over a time period you specify, rather than sending a full proportion of the traffic immediately. Slow Ramp is essential when using the Least Connections load balancing method (our recommended method), as the BIG-IP system would otherwise send all new connections to a new server immediately, potentially overwhelming that server. Select whether you want to use a Slow Ramp time.

- **Use Slow Ramp**
Select this option for the system to implement Slow Ramp time for this pool.
 - How many seconds should Slow Ramp time last?***
Specify a duration in seconds, for Slow Ramp. The time period you select for Slow Ramp is highly dependent on the speed of your server hardware and the behavior of your web services. The default setting of 300 seconds (5 minutes) is very conservative in most cases.
- **Do not use Slow Ramp**
Select this option if you do not want to use Slow Ramp. If you select this option, we recommend you do not use the Least Connections load balancing method.

Application Health

In this section, you answer questions about how you want to implement application health monitoring on the BIG-IP system.

1. ***Create a new health monitor or use an existing one?***

Application health monitors are used to verify the content that is returned by an HTTP request. The system uses these monitors to ensure traffic is only sent to available servers.


Unless you have requirements for configuring other options not in the following list of questions, we recommend allowing the iApp to create a new monitor. Creating a custom health monitor is not a part of this template; see **Local Traffic >> Monitors**. To select any new monitors you create, you need to restart or reconfigure this template.

- **Select the monitor you created from the list**
If you manually created the health monitor, select it from the list.
Continue with *iRules on page 22*.
- **Create a new health monitor**
If you want the iApp to create a new monitor, continue with the following.
 - a. **How many seconds should pass between health checks?**
Specify how long the system should wait between each health check. This is used as the Interval setting for the monitor.
We recommend the default of 30 seconds.
 - b. **What type of HTTP request should be sent to the servers?**
Select whether you want the system to send an HTTP GET or POST request. The GET method requests data from the server, the POST submits data to be processed by the server.
 - **GET**
Select this option if you want the system to use a GET request. The system uses the URI you specify in the next question to request content from the server.
 - **POST**
Select this option if you want the system to use a POST request. The system uses the URI you specify in the next question, along with the HTTP POST body you will specify to form the request.
 - c. **What HTTP URI should be sent to the servers?**
The HTTP URI is used to specify the resource on the Office Web Apps server for a given request. This parameter can be customized to request a specific part of your application, which can indicate the application-health on a granular level. By default, the URI for Office Web Apps is **/hosting/discovery**.
 - d. **What HTTP version do your servers expect clients to use?**
Choose the HTTP version which you expect most clients are using. This allows the system to detect failures more accurately.
 - **HTTP/1.0**
Choose this option if you expect your clients to use HTTP/1.0.
 - **HTTP/1.1**
Choose this option if you expect your clients to use HTTP/1.1.
 - e. **What HTTP POST body do you want to use for this monitor?**
This question only appears if you selected a POST request.

If you selected a POST request, you must specify the message body for the POST.
 - f. **What is the expected response to the HTTP request?**
Specify the response you expect returned from the request. The system checks the response from the server against the response you enter here to determine server health. By default, the response for Office Web Apps is **wopi-discovery**.
 - g. **Should the health monitor require credentials?**
Choose whether the system should try to authenticate to the Office Web Apps deployment as part of the health check.
 - **No, allow anonymous access**
Select this option if you do not want the monitor to attempt authentication.
 - **Yes, require credentials for Basic authentication**
Select this option if you want to attempt Basic authentication as a part of the health monitor. To require credentials, you should have a user account specifically for this health monitor which has no other privileges, and has a password set to never expire.
 - a. **What user name should the monitor use?**
Type the domain and user name for the account you created for the health monitor. You must include the domain in front of the user, such as EXAMPLE\USER.
 - b. **What is the associated password?**
Type the password for the account.
 - **Yes, require credentials for NTLM authentication**
Select this option if you want to attempt NTLM authentication as a part of the health monitor. To require credentials, you should have a user account specifically for this monitor with no other privileges, and a password set to never expire.

a. What user name should the monitor use?

Type the user name for the account you created for the health monitor.

 **Important** *If you have a complex Active Directory environment where the username is not in the same domain as the Domain Controller being queried, you should prepend the domain name in uppercase letters (DOMAIN\) to the username in this field.*

b. What is the associated password?


Type the password for the account.

iRules

In this section, you can add custom iRules to the deployment. This entire section is available only if you selected Advanced mode. iRules are a scripting language that allows an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. An iRule contains the set of instructions the system uses to process data flowing through it, either in the header or payload of a packet.

1. Do you want to add any custom iRules to the configuration? **Advanced**

Select if you have preexisting iRules you want to add to your Office Web Apps implementation.

 **Warning** *While iRules can provide additional functionality not present in the iApp, improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommend you verify the impact of an iRule prior to deployment in a production environment.*

If you do not want to add any iRules to the configuration, continue with the following section.


If you have iRules you want to attach to the virtual server the iApp creates for your Office Web Apps servers, from the **Options** box, click the name of the applicable iRule(s) and then click the Add (<<) button to move them to the **Selected** box.

Statistics and Logging

This section contains questions about optional logging and statistics. This section is available only if you selected Advanced mode.

1. Do you want to enable Analytics for application statistics?

The Application Visibility Reporting (AVR) module for analytics allows you to view statistics specific to your application implementation. AVR is included and available on all BIG-IP systems v11 and later, however you must have the AVR provisioned for this option to appear. Note that this provisioning requirement is only for AVR, you can view object-level statistics from the BIG-IP system without provisioning AVR.

 **Important** *Enabling Analytics may adversely affect overall system performance. If you choose to enable Analytics, we recommend gathering statistics for a set time period, such as one week, and then re-entering this template and disabling Analytics while you process the data.*

If you plan on using AVR for analytics, we recommend creating a custom Analytics profile. To create a new profile, from the Main tab, select **Profiles** and then click **Analytics**. Click **New** and then configure the profile as applicable for your configuration. See the online help or product documentation for specific instructions. To select the new profile, you need to restart or reconfigure the iApp.

- **Do not enable Application Visibility Reporting**

If you do not want to enable Analytics, leave this list set to **No**, and continue with the next section.

- **Select the Analytics profile you created from the list**

If you choose to enable Analytics, select the Analytics profile you want to use for this implementation from the list.

2. Which HTTP request logging profile do you want to use?

HTTP request logging enables customizable log messages to be sent to a syslog server for each HTTP request processed by your application. You can choose to enable HTTP request logging by selecting a logging profile you already created from the list.

We strongly recommend you thoroughly test the performance impact of using this feature in a staging environment prior to enabling on a production deployment

Creating a request logging profile is not a part of this template. See Local Traffic>>Profiles: Other: Request Logging. To select any new profiles you create, you need to restart or reconfigure this template.

- **Do not enable HTTP request logging**
If you do not want to enable HTTP request logging, leave this list set to **No**, and continue with the next section.
- **Select the HTTP request logging profile you created from the list**
If you choose to enable HTTP request logging, select the profile you want to use for this implementation from the list.

Finished

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

Archived

Next steps

After completing the iApp Template, the BIG-IP Application Services page opens for the Office Web Apps service you just created. To see the list of all the configuration objects created to support the Office Web Apps application, on the Menu bar, click **Components**. The complete list of all 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 Office Web Apps 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 Office Web Apps 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 configuration objects created by the iApp template. You can get statistics specific to the Application Service if you have provisioned AVR. Otherwise, you can always get object-level statistics.

AVR statistics

If you have provisioned AVR, you can get application-level statistics for your Office Web Apps Application Service.

To view AVR statistics

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. From the Application Service List, click the Office Web Apps service you just created.
3. On the Menu bar, click **Analytics**.
4. Use the tabs and the Menu bar to view different statistics for your iApp.

Object-level statistics

If you haven't provisioned AVR, or want to view object-level statistics, use the following procedure.

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

We strongly recommend using the iApp template to configure the BIG-IP system for Office Web Apps. Users familiar with the BIG-IP system can use the following table to manually configure the BIG-IP system. The table contains a list of configuration objects along with any non-default settings you should configure as a part of this deployment. Settings not mentioned in the table can be configured as applicable for your configuration. For specific instructions on configuring individual objects, see the online help or product manuals.

Health Monitors (Local Traffic > Monitors)		
Name	Type a unique name	
Type	HTTP (HTTPS if you are deploying SSL Bridging)	
Interval	30 (recommended)	
Timeout	91 (recommended)	
Send String	GET /hosting/discovery HTTP/1.1\r\nHost: wac.example.com\r\nConnection: Close\r\n\r\n Replace the red text with your FQDN.	
Receive String	wopi-discovery	
User Name¹	If necessary, type a user name with access to your application. We recommend creating an account for use in this monitor. If you are using NTLM authentication and have a complex Active Directory environment where the username is not in the same domain as the Domain Controller being queried, you must prepend the domain name in uppercase letters (DOMAIN) to the username in this field.	
Password¹	Type the associated password	
Pools (Local Traffic > Pools)		
Name	Type a unique name	
Health Monitor	Select the monitor you created above	
Slow Ramp Time¹	300 (recommended)	
Load Balancing Method	Choose a load balancing method. We recommend Least Connections (Member)	
Address	Type the IP Address of the nodes	
Service Port	80 (if performing SSL offload, otherwise use 443 . Click Add to repeat Address and Service Port for all nodes.	
Optional: AAM Application ² (Acceleration > Web Application)		
Application Name	Type a unique name	
Policy	Generic Policy - Enhanced	
Requested Host	Type the Fully Qualified Domain Name (FQDN) of your application. Click Add Host to include additional hosts.	
Profiles (Local Traffic > Profiles)		
HTTP (Profiles > Services)	Name	Type a unique name
	Parent Profile	http
	Rewrite Redirect ²	Matching
TCP WAN (Profiles > Protocol)	Name	Type a unique name
	Parent Profile	tcp-wan-optimized
TCP LAN (Profiles > Protocol)	Name	Type a unique name
	Parent Profile	tcp-lan-optimized
Persistence (Profiles > Persistence)	Name	Type a unique name
	Persistence Type	Cookie
OneConnect (Profiles > Other)	Name	Type a unique name
	Parent Profile	oneconnect
Client SSL³ (Profiles > SSL)	Name	Type a unique name
	Parent Profile	clientssl
	Certificate and Key	Select the Certificate and Key you imported from the associated list
Server SSL⁴ (Profiles > Other)	Name	Type a unique name
	Parent Profile	serverssl
Web Acceleration (Profiles > Services)	Name	Type a unique name
	Parent Profile	optimized-caching
	WA Applications ⁵	Enable the AAM Application you created

¹ User name and password are optional and only required if the health monitor should require credentials.

² You must select **Advanced** from the **Configuration** list for these options to appear

³ Optional. The BIG-IP AAM configuration is recommended, but optional.

⁴ Only required if using the BIG-IP system for SSL Offload or SSL Bridging

⁵ Only necessary if using the BIG-IP system for SSL Bridging or server-side encryption

Profiles continued		
iSession⁵ (Profiles > Services)	Name Parent Profile	Type a unique name isession
HTTP Compression (Profiles > Services)	Name Parent Profile	Type a unique name wan-optimized-compression
	Content List--> Include List (Add each entry to the Content Type box and then click Include)	text/(css html javascript json plain postscript richtext rtf vnd .wap .wml vnd .wap .wmlscript wap wml x-component x-vcalendar x-vcard+xml)
		application/(css css-stYLESHEET doc excel javascript json lotus123 mdb mpp ms-excel ms-powerpoint ms-word msaccess msexcel mspoverpoint msproject msword photoshop postscript powerpoint ps psd quarkexpress rtf txt visio vnd .excel vnd .ms-access vnd .ms-excel vnd .ms-powerpoint vnd .ms-pps vnd .ms-project vnd .ms-word vnd .ms-works vnd .ms-works-db vnd .msaccess vnd .msexcel vnd .mspoverpoint vnd .msword vnd .powerpoint vnd .visio vnd .wap .cmlscriptc vnd .wap .wmlc vnd .wap .xhtml .xhtml .xhtml .xhtml .word vnd .winword wks word x-excel x-java-jnlp-file x-javascript x-json x-lotus123 x-mdb x-msexcel x-ms-project x-mscardfile x-msclip x-msexcel x-mspowerpoint x-msproject x-msword x-msworks-db x-msworks-wps x-photoshop x-postscript x-powerpoint x-ps x-quark-express x-rtf x-vermeer-rpc x-visio x-vsd x-wks x-word x-xls x-xml xhtml+xml xls+xml)
	image/(photoshop psd x-photoshop x-vsd)	
Virtual Servers (Local Traffic > Virtual Servers)		
HTTP		
Name	Type a unique name.	
Address	Type the IP Address for the virtual server	
Service Port	80	
Protocol Profile (client)^{1,2}	Select the WAN optimized TCP profile you created	
Protocol Profile (server)^{1,2}	Select the LAN optimized TCP profile you created	
HTTP Profile²	Select the HTTP profile you created	
Web Acceleration profile²	Select the Web Acceleration profile you created	
HTTP Compression profile²	Select the HTTP Compression profile you created	
OneConnect²	Select the OneConnect profile you created	
Source Address Translation³	Auto Map (optional; see footnote ³)	
iSession profile⁵	If using BIG-IP AAM for symmetric optimization between systems, select the iSession profile you created.	
Default Pool²	Select the pool you created	
Persistence Profile²	Select the Persistence profile you created	
iRule⁴	If offloading SSL only: Enable the built-in _sys_https_redirect irule	
HTTPS ⁴		
Name	Type a unique name.	
Address	Type the IP Address for the virtual server	
Service Port	443	
Protocol Profile (client)¹	Select the WAN optimized TCP profile you created	
Protocol Profile (server)¹	Select the LAN optimized TCP profile you created	
HTTP Profile	Select the HTTP profile you created above	
Web Acceleration profile	Select the Web Acceleration profile you created	
HTTP Compression profile	Select the HTTP Compression profile you created	
OneConnect	Select the OneConnect profile you created	
SSL Profile (Client)	Select the Client SSL profile you created	
SSL Profile (Server)⁶	If you created a Server SSL profile, select it from the list	
Source Address Translation³	Auto Map (optional; see footnote ³)	
iSession profile⁵	If using BIG-IP AAM for symmetric optimization between systems, select the iSession profile you created.	
Default Pool	Select the pool you created	
Persistence Profile	Select the Persistence profile you created	

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

² Do not enable these objects on the HTTP virtual server if offloading SSL. The HTTP virtual server is only used for redirecting users to the HTTPS virtual server.

³ If using SNAT and expecting more than 64,000 simultaneous connections, you must configure a SNAT Pool with an IP address for each 64,000 simultaneous connections you expect. See the BIG-IP documentation on configuring SNAT Pools.

⁴ Only necessary if offloading SSL or SSL Bridging

⁵ Only necessary if using the BIG-IP AAM to provide symmetric optimization. Do not create/use this profile if you are deploying the BIG-IP on the server side of the WAN

⁶ Only necessary if using the BIG-IP system for SSL Bridging or server-side encryption

Adding Office Web Apps support to a SharePoint 2013 virtual server

If you have an existing BIG-IP deployment for SharePoint 2013, you can add support for Office Web Apps by adding an iRule to the existing virtual server.

If you are combining Office Web Apps with the SharePoint virtual server, you must use a Subject Alternative Name certificate that contains the host names of both the SharePoint and Office Web Apps farms.

If you are using BIG-IP AAM (formerly WebAccelerator) on your existing SharePoint 2013 deployment, you must add the Office Web Apps host name to the Application in the Requested Hosts field. For specific instructions, see the BIG-IP AAM documentation.

Creating the health monitor and pool for the Office Web Apps servers

Use the following table for guidance on configuring the BIG-IP LTM for Office Web Apps. For specific instructions on configuring these objects, see the online help or the BIG-IP documentation.

Health Monitors (Local Traffic > Monitors)		
Name	Type a unique name	
Type	http (if performing SSL Offload - recommended), otherwise use https	
Interval	30 (recommended)	
Timeout	91 (recommended)	
Send String	GET /hosting/discovery HTTP/1.1\r\nHost: wac.example.com\r\nConnection: Close\r\n\r\n Replace the red text with your FQDN.	
Receive String	wopi-discovery	
Pools (Local Traffic > Pools)		
Pool (Main tab-->Local Traffic -->Pools)	Name	Type a unique name, such as office-web-apps-pool
	Health Monitor	Select the monitor you created above
	Slow Ramp Time	300
	Load Balancing Method	Least Connections (Member)
	Address	Type the IP Address of an Office Web Apps server
	Service Port	80 (if performing SSL Offload - recommended), otherwise use 443 Click Add to repeat Address and Service Port for all nodes

Creating the iRule

The first task is to create the iRule. This iRule sends Office Web Apps traffic to the correct pool of servers.

To create the iRule

1. On the Main tab, expand **Local Traffic**, and then click **iRules**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this iRule.
4. In the **Definition** section, copy and paste the following iRule, omitting the line numbers.

```
1  when HTTP_REQUEST {
2      if { [string tolower [HTTP::host]] contains "wac.example.com" } {
3          pool my_officewebapps_pool
4      } else {
5          pool my_sharepoint_2013_pool
6          persist none
7      }
8  }
```

5. Click **Finished**.

Adding the iRule to the SharePoint 2013 virtual server

The next task is to modify the BIG-IP virtual server for SharePoint 2013 to use the iRule you just created. This could be a virtual server created by the SharePoint iApp template, or created manually.

Disabling the Strict Updates feature (if you used the iApp template only)

If you used the iApp template to configure the BIG-IP system for SharePoint 2013, before modifying the virtual server, you must turn off the Strict Updates feature. By turning off Strict Updates, if you re-enter the iApp template and modify the configuration within the iApp, you will have to make all of the following changes again manually.

To turn off Strict Updates

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. Click the name of your SharePoint Application service from the list.
3. From the **Application Service** list, select **Advanced**.
4. In the **Strict Updates** row, clear the check from the box to disable Strict Updates.
5. Click the **Update** button.

Adding the iRule to the virtual server

The final task is to modify the SharePoint HTTPS virtual server to use the iRule you created.

To add the iRule to the SharePoint 2013 virtual server

1. On the Main tab, expand **Local Traffic** and then click **Virtual Servers**.
2. From the list, click the name of the SharePoint HTTPS virtual server created by the iApp or manually. This virtual server is preceded by the name you gave the iApp, followed by **_https_virtual**.
3. On the Menu bar, click **Resources**.
4. In the iRules section, click **Manage**.
5. From the **Available** list, click the name of the iRule you created, and then click the Add (<<) button to enable it.
6. Click **Finished**.

BIG-IP Access Policy considerations for Office Web Apps server

This section includes guidance on modifying the BIG-IP system configuration to support clients accessing Office Web Apps through a SharePoint 2013 virtual server that is secured with F5's Access Policy Manager (APM) module.

Office Web Apps server(s) must be able to retrieve files from the SharePoint server(s) prior to rendering them in a browser. If you have secured authentication to SharePoint 2013 using F5's APM module, and the Office Web Apps host name resolves to the IP address of the SharePoint 2013 virtual server, you must create an internal BIG-IP virtual server to allow requests for content directly from Office Web Apps servers without pre-authentication. You must also configure host file entries on the Office Web Apps server(s) to forward requests to this internal BIG-IP virtual server (see the Microsoft documentation for specific details).

Also, because clients access Office Web Apps from the same client session using a different fully qualified domain name, the BIG-IP APM Access Policy must be configured to support multiple domain names, granting access to the Office Web Apps farm FQDN to the previously authenticated client APM session for SharePoint 2013 or Exchange 2013."

Adding multiple host domains to the Access Profile

In this section, we modify the existing Access Profile for SharePoint or Exchange Server to add multiple host domains.

Disabling Strict Updates

If your SharePoint or Exchange Server configuration was created using an iApp template, you may have to disable the Strict Updates feature first. If you did not use an iApp template, continue with the following section.

To disable Strict Updates

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. Click the name the Application service for your Microsoft application from the list.
3. If necessary, from the **Application Service** list, select **Advanced**.
4. In the **Strict Updates** row, clear the check box to disable Strict Updates.
5. Click **Update**.

Modifying the Access Profile

Use the following procedure to modify the Access Profile for the Microsoft application.

To modify the Access Profile

1. On the Main tab, expand **Access Policy** and then click **Access Profiles**.
2. Click the Access Profile used by the BIG-IP virtual server for your Microsoft application.
3. On the Menu bar, click **SSO/Auth Domains**.
4. In the **Domain Cookie** box, if there is any cookie specified, clear the box and then click the **Update** button.
5. In the **Domain Mode** row, click the **Multiple Domains** button.
6. In the **Primary Authentication URI** box, type the full path of the site where users authenticate, such as *https://sp.example.com*.
7. Verify SSO Configuration setting is correct.
8. In the **Cookie** row, select **Domain** from the list, and then type the primary domain name to which users will be authenticating.
9. *SharePoint 2013 deployments only*: In the **Cookie Options** row, click the **Persistent** box.
10. Click **Update**.
11. In the Authentication Domains section, click the **Add** button on the right.
12. Click **Add**. Add another domain corresponding to the FQDN of the Office Web Apps farm.
13. In the **Cookie** row, select **Domain** from the list, and then type the primary domain name to which users will be authenticating.
14. From the **SSO Configuration** list, select the same SSO Configuration as in step 7.
15. Click **Update**.
16. In the upper left corner of the screen, click the **Apply Access Policy** link.
You should now be able to authenticate once for both the primary application and Office Web Apps domains.

Creating the internal virtual server on the BIG-IP system

Use the following table for guidance on configuring internal virtual server(s) on the BIG-IP system for the Microsoft applications. This table contains a list of BIG-IP LTM configuration objects along with any non-default settings you should configure as a part of this deployment. For specific instructions on configuring individual objects, see the online help or product manuals.

If you are deploying Office Web Apps with Microsoft Exchange 2013, the pool members for the internal virtual server only need to be the Outlook Web Access servers.

Health Monitor (<i>Local Traffic > Health Monitor</i>)	
Name	Type a unique name
<i>Note: This settings for this health monitor should match the corresponding health monitor for the Microsoft application you created (or that was created by the iApp). The settings in the following example are for a simple, generic health monitor.</i>	
Type	http (if performing SSL Offload), otherwise use https
Interval	30 (recommended)
Timeout	91 (recommended)

Pools (*Local Traffic > Pools*)

Name	Type a unique name, such as exchange-internal-pool
Health Monitor	Select the monitor you created above
Slow Ramp Time¹	300
Load Balancing Method	Least Connections (Member)
Address	Type the IP Address of an appropriate Microsoft application server.
Service Port	80 or 443 , depending on the configuration of the application servers. Click Add to repeat Address and Service Port for all nodes

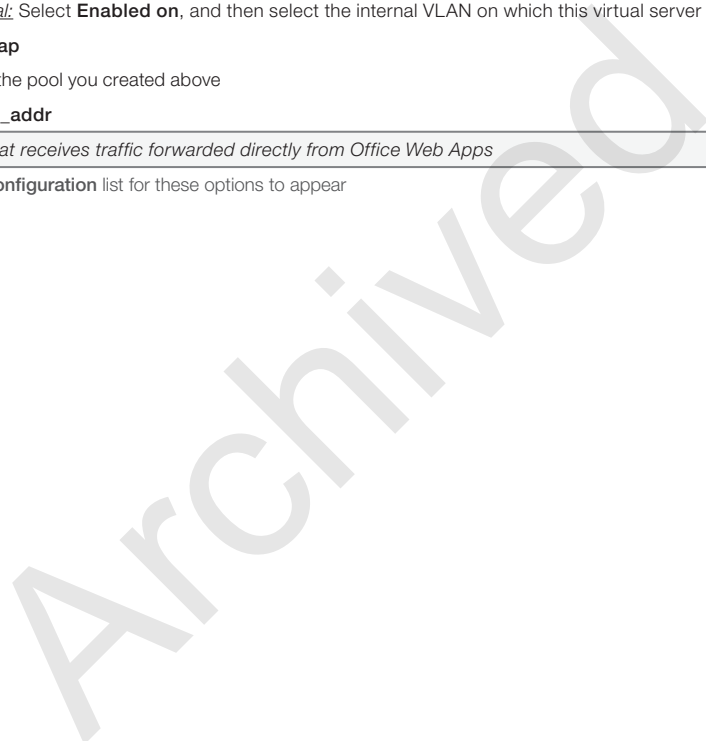
Virtual Servers (*Local Traffic > Virtual Servers*)

Name	Type a unique name.
Address	IP address to which DNS requests for the Microsoft application will point. This should be on the same internal network as the Web Apps server.
Service Port	80 or 443 , depending on your configuration
Protocol Profile (Client)¹	tcp-ian-optimized
VLAN and Tunnel Traffic	<i>Optional:</i> Select Enabled on , and then select the internal VLAN on which this virtual server should listen.
Source Address Translation	Automap
Default Pool	Select the pool you created above
Persistence Profile	source_addr

Repeat this table for each application that receives traffic forwarded directly from Office Web Apps

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

This completes the configuration.



Troubleshooting

This section contains information on troubleshooting common problems.

Q: *How do I allow mixed content for Office Web Apps connections?*

A: When using Office Web Apps Server 2013 to preview documents in SharePoint 2013, Exchange 2013 Outlook Web App, or the Lync 2013 client, you may be prompted repeatedly by Internet Explorer to allow viewing of mixed content. In the case of the Lync 2013 client, the sharing of PowerPoint presentations may hang or fail. To remediate this behavior, add the URL of the Office Web Apps farm to **Trusted Sites** in Internet Explorer options, and ensure that the **Enable** radio button is selected in the **Display Mixed Content** section of the **Security Settings - Trusted Sites Zone**.

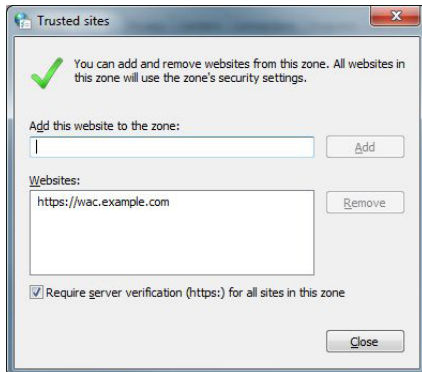


Figure 1: *Trusted Sites configuration in Internet Explorer*

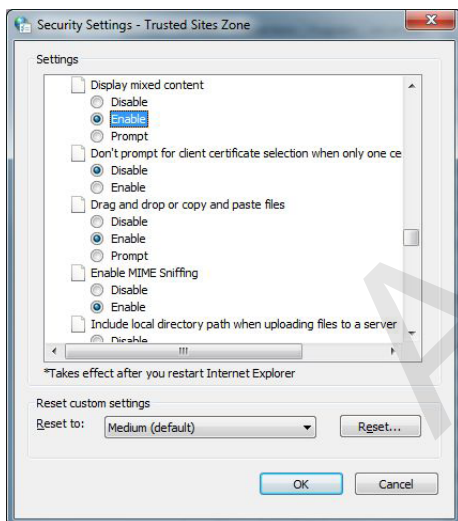


Figure 2: *Trusted Sites configuration in Internet Explorer*

For specific information on configuring Internet Explorer, see the Microsoft documentation.

Q: *Why are client connections unresponsive or seem to hang when using the OneConnect feature?*

A: If you have configured your deployment to use OneConnect (part of F5's recommended configuration), and users are experiencing slow performance or the need to refresh pages, Microsoft IIS may be failing to reset the TCP connection after the default timeout period of 120 seconds.

To work around this issue, create a custom server-side TCP profile with an **Idle Timeout** value of less than 120 seconds, and then apply the profile to the virtual server. If you used the iApp template to configure the BIG-IP system, you can attach the new profile using the template. If you manually configured the BIG-IP system, you manually add the profile.

To create the new TCP profile, click **Local Traffic > Profiles > Protocol > TCP** and then click **Create**. From the Parent Profile list, select **tcp-lan-optimized**. Check the Custom box for **Idle Timeout**, and then in the **Seconds** box, type a number less than 120, such as **110**. Click **Finished** to create the profile.

To add the profile to the virtual server

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. Click the name of your Office Web App Application service from the list.
3. On the Menu bar, click **Reconfigure**.
4. If necessary, from the *Which configuration mode do you want to use?* question, select **Advanced - Configure advanced options**.
5. In the Server Offload section, from the *Which OneConnect profile do you want to use?* question, select the profile you just created.
6. Click the **Finished** button.

If you configured the BIG-IP system manually, simply replace the existing **Protocol Profile (Server)** profile with the profile you just created.

Q: *Why are Office 2011 clients unable to view or edit Office documents?*

A: If you have collocated Office Web Apps and SharePoint behind a BIG-IP APM virtual server (see page 27), Macintosh clients running Office 2011 may be unable to view or edit Office documents. You must disable APM policy processing for requests from Macintosh clients by attaching the following iRule to the SharePoint/Office Web Apps BIG-IP virtual server (you may simply replace the rule you configured on the bottom of page 27 with the following). Be sure to change the text in lines 2, 6, and 8 to match your configuration.

```
1  when HTTP_REQUEST {
2      if { [string tolower [HTTP::host]] contains "office-web-apps-fqdn" } {
3          if { [string tolower [HTTP::header "User-Agent"]] contains "macintosh" } {
4              ACCESS::disable
5          }
6          pool my_officewebapps_pool
7      } else {
8          pool my_sharepoint_2013_pool
9          persist none
        }
    }
```

➡ **Note:** *You must also have created the internal BIG-IP virtual server for receiving requests for SharePoint content from the Office Web Apps servers, as described in [Creating the internal virtual server on the BIG-IP system on page 29](#).*

Glossary

application service

iApps Application Services use an [iApp Template](#) to guide users through configuring new BIG-IP® system configurations. An Application Service lets an authorized user easily and consistently deploy complex BIG-IP® system configurations just by completing the information required by the associated template. Every Application Service is attached to a specific configuration and cannot be copied the way that iApps templates can.

iApp Template

iApps templates create configuration-specific forms used by Application Services to guide authorized users through complex system configurations. The templates provide programmatic, visual layout and help information. Each new Application Service uses one of the templates to create a screen with fields and help that guide the user through the configuration process and creates the configuration when finished.

iApps templates allow users to customize by either modifying an existing template or creating one from scratch. Users can create scratch-built templates using either the iApps Templates screen or any text-editing software.

configuration utility

The Configuration utility is the browser-based application that you use to configure the BIG-IP system.

custom profile

A custom [profile](#) is a profile that you create. A custom profile can inherit its default settings from a parent profile that you specify. See also parent profile.

health monitor

A health monitor checks a node to see if it is up and functioning for a given service. If the node fails the check, it is marked down. Different monitors exist for checking different services.

iRule

An iRule is a user-written script that controls the behavior of a connection passing through the BIG-IP system. iRules™ are an F5 Networks feature and are frequently used to direct certain connections to a non-default load balancing pool. However, iRules can perform other tasks, such as implementing secure network address translation and enabling session persistence. You can attach iRules you created to your HTTP Application Service in the advanced configuration mode.

iSession

An iSession is an optimized connection between two BIG-IP systems.

iSession profile

An iSession profile defines the optimization parameters. WAN optimization requires an iSession profile, which specifies the optimization settings, such as compression and data deduplication. The iApp template uses the default isession profile.

load balancing method

A load balancing method or algorithm is a particular method of determining how to distribute connections across a [load balancing pool](#). There are several different load balancing methods on the BIG-IP system. If you are working with servers that differ significantly in processing speed and memory, you might want to use a method such as Ratio or Weighted Least Connections.

Load balancing calculations can be localized to each pool (member-based calculation) or they may apply to all pools of which a server is a member (node-based calculation). For detailed information, see the product documentation.

See the table on the following page for a description of most load balancing methods.

Method	Description	When to use
Round Robin	Round Robin mode passes each new connection request to the next server in line, eventually distributing connections evenly across the array of machines being load balanced.	Round Robin mode works well in most configurations, especially if the equipment that you are load balancing is roughly equal in processing speed and memory.
Ratio (member) Ratio (node)	The LTM distributes connections among pool members in a static rotation according to ratio weights you define. The number of connections each system receives over time is proportionate to the ratio weight you defined for each pool member. You set a ratio weight when you add each pool member in the iApp.	These are static load balancing methods, basing distribution on user-specified ratio weights that are proportional to the capacity of the servers.
Dynamic Ratio (member) Dynamic Ratio (node)	The Dynamic Ratio methods select a server based on various aspects of real-time server performance analysis. These methods are similar to the Ratio methods, except the ratio weights are system-generated, and the values of the ratio weights are not static. These methods are based on continuous monitoring of the servers, and the ratio weights are therefore continually changing.	The Dynamic Ratio methods are used specifically for load balancing traffic to RealNetworks® RealSystem® Server platforms, Windows® platforms equipped with Windows Management Instrumentation (WMI), or any server equipped with an SNMP agent such as the UC Davis SNMP agent or Windows 2000 Server SNMP agent. Note: To implement Dynamic Ratio load balancing, you must first install and configure the necessary server software for these systems, and then install the appropriate performance monitor.
Fastest (node) Fastest (application)	The Fastest load balancing mode load balances based upon the number of outstanding Layer 7 requests to a pool member and the number of open L4 connections.	The Fastest methods are useful in environments where nodes are distributed across separate logical networks.
Least Connections (member) Least Connections (node)	The Least Connections load balancing mode is a dynamic load balancing algorithm that distributes connections to the server that is currently managing the fewest open connections at the time the new connection request is received.	The Least Connections methods function best in environments where the servers have similar capabilities. Otherwise, some amount of latency can occur. If you have servers with varying capacities, consider using the Weighted Least Connections methods instead.
Weighted Least Connections (member) Weighted Least Connections (node)	Specifies that the system passes a new connection to the pool member that is handling the lowest percentage of the specified maximum number of concurrent connections allowed. This mode requires that you specify a value for the connection-limit option for all members of the pool.	This mode works best in environments where the servers or other equipment you are load balancing have different but quantified capability limits.
Observed (member) Observed (node)	With the Observed methods, nodes are ranked based on the number of connections. The Observed methods track the number of Layer 4 connections to each node over time and create a ratio for load balancing	The need for the Observed methods is rare, and they are not recommended for large pools.
Predictive (member) Predictive (node)	The Predictive methods use the ranking methods used by the Observed methods. However, with the Predictive methods, LTM analyzes the trend of the ranking over time, determining whether a nodes performance is currently improving or declining. The servers with performance rankings that are currently improving receive a higher proportion of the connections.	The need for the Predictive methods is rare, and they are not recommended for large pools.
Least Sessions	The Least Sessions method selects the server that currently has the least number of entries in the persistence table. Use of this load balancing method requires that the virtual server reference a type of profile that tracks persistence connections, such as the Source Address Affinity or Universal profile type. Note: The Least Sessions methods are incompatible with cookie persistence.	The Least Sessions method works best in environments where the servers or other equipment that you are load balancing have similar capabilities.

load balancing pool

A load balancing pool is a logical set of devices, such as web servers, that you group together to receive and process traffic. Instead of sending client traffic to the destination IP address specified in the client request, Local Traffic Manager sends the request to any of the servers that are members of that pool. This helps to efficiently distribute the load on your server resources.

local endpoint

The local endpoint is the BIG-IP system on which you are currently working. The systems must be set up symmetrically, so that a local endpoint connects to one or more remote endpoints.

network virtual server

A network virtual server is a virtual server whose IP address has no bits set in the host portion of the IP address (that is, the host portion of its IP address is 0, such as 192.168.1.0). This allows you to direct client traffic based on a range of destination IP addresses.

profile

Profiles are a configuration tool that you can use to affect the behavior of certain types of network traffic. More specifically, a profile is an object that contains settings with values, for controlling the behavior of a particular type of network traffic, such as HTTP connections. Profiles also provide a way for you to enable connection and session persistence, and to manage client application authentication.

self IP address

Self IP addresses are the IP addresses owned by the BIG-IP system that you use to access the internal and external VLANs.

SNAT

A SNAT (Secure Network Address Translation) is a feature that defines a routable alias IP address that one or more nodes can use as a source IP address when making connections to hosts on the external network.

SNAT pool

A SNAT pool is a pool of translation addresses that you can map to one or more original IP addresses. Translation addresses in a SNAT pool are not self IP addresses.

virtual server

A virtual server is a traffic-management object on the BIG-IP system that is represented by an IP address and a service port. This is the address clients use to connect to the web servers (or a FQDN resolves to this address). The BIG-IP intercepts the client request, and then directs the traffic according to your configuration instructions.

VLAN

A VLAN is a logical grouping of interfaces connected to network devices. You can use a VLAN to logically group devices that are on different network segments. Devices within a VLAN use Layer 2 networking to communicate and define a broadcast domain.

Document Revision History

Version	Description	Date
1.0	New document	12-11-2012
1.1	Added the Troubleshooting section page 31, with an entry on Allowing Mixed Content for Office Web Apps Connections.	01-11-2013
2.0	Updated guide for the downloadable iApp template.	01-15-2014
2.1	Added support for BIG-IP v11.4.1 and 11.5.	01-31-2014
2.2	Added an entry to the Troubleshooting section page 31 for client connections that are unresponsive or seem to hang when using the OneConnect feature.	06-20-2014
2.3	Added support for BIG-IP v11.5.1 and 11.6.	08-25-2014
2.4	Added a Troubleshooting entry page 32, concerning Office 2011 clients being unable to view or edit office documents.	12-18-2014
2.5	Added support for BIG-IP v12.0	09-02-2015
2.6	Added a note to health monitor section in the iApp walkthrough concerning when to prepend Domain information if using NTLM authentication. Added the same note to the manual configuration table.	05-02-2016
2.7	Added support for BIG-IP v12.1	05-18-2016

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

