



## What's inside:

- 2 What is F5 iApp™?
- 2 Prerequisites and configuration notes
- 5 Deployment Scenarios
- 7 Preparation worksheets
- 9 Downloading and importing the new iApp template
- 9 Configuring the BIG-IP iApp for Microsoft Exchange Server 2010
- 35 Modifying the iApp template configuration
- 42 Next steps
- 43 Appendix A: Configuring DNS and NTP settings
- 44 Appendix B: Using X-Forwarded-For to log the client IP address in IIS 7.0 & 7.5
- 46 Appendix C: Manual configuration tables
- 69 Appendix D: Technical Notes
- 72 Revision History



## Deploying the BIG-IP v11 with Microsoft Exchange 2010 Client Access Servers

Welcome to the F5 and Microsoft® Exchange® 2010 Client Access Server deployment guide. This document contains guidance on configuring the BIG-IP system version 11 and later for the Client Access Service of Exchange 2010, including SP1 and SP2, resulting in a secure, fast, and available deployment. BIG-IP version 11.0 introduces iApp™ Application templates, an extremely easy and accurate way to configure the BIG-IP system for Microsoft Exchange Server 2010.

This document provides guidance for using the **new, downloadable BIG-IP iApp** Template to configure the Client Access server role of Microsoft Exchange Server 2010, as well as instructions on how to configure the BIG-IP system manually.

By using the iApp template, you can configure the BIG-IP system to support any combination of the following services supported by Client Access servers: Outlook Web App (which includes the HTTP resources for Exchange Control Panel, Exchange Web Services, and Offline Address Book), Outlook Anywhere (RPC over HTTP), ActiveSync, Autodiscover, RPC Client Access (MAPI), POP3 and IMAP4. This guide also contains manual configuration instructions for users familiar with F5 devices.

For more information on the Client Access Server role, see [technet.microsoft.com/en-us/library/bb124915%28EXCHG.140%29.aspx](http://technet.microsoft.com/en-us/library/bb124915%28EXCHG.140%29.aspx)

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

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

### Why F5?

F5 offers a complete suite of application delivery technologies designed to provide a highly scalable, secure, and responsive Exchange 2010 deployment.

- Terminating HTTPS connections at the BIG-IP LTM reduces CPU and memory load on Client Access Servers, and simplifies TLS/SSL certificate management.
- The BIG-IP LTM can balance load and ensure high-availability across multiple Client Access servers using a variety of load-balancing methods and priority rules.
- The BIG-IP LTM's TCP Express feature set ensures optimal network performance for all clients and servers, regardless of operating system and version.
- The LTM provides content compression features which improve client performance.
- The BIG-IP Access Policy Manager or Edge Gateway, F5's high-performance access and security solutions, can provide proxy authentication and secure remote access to Exchange HTTP-based Client Access services.

**Important:** Make sure you are using the most recent version of this deployment guide, available at <http://www.f5.com/pdf/deployment-guides/microsoft-exchange2010-iapp-dg.pdf>

### What is F5 iApp™?

New to BIG-IP version 11, F5 iApp is a powerful new set of features in the BIG-IP system that provides a new way to architect application delivery in the data center. iApp 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 Exchange Server 2010 acts as the single-point interface for building, managing, and monitoring the Exchange 2010 Client Access role.

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 prerequisites and configuration notes for the Client Access Role:

**Critical**



- This document provides guidance on using the **downloadable iApp** for Microsoft Exchange 2010 found at <http://support.f5.com/kb/en-us/solutions/public/13000/400/sol13497.html>, and **not** the iApp found by default in BIG-IP version 11. **You must use this downloadable iApp for BIG-IP versions 11.0, 11.0.1 and 11.1**, as it contains a number of fixes and enhancements not found in the default iApp. Future versions of the BIG-IP system will include this iApp by default. You must have a current support contract and associated account on downloads.f5.com to download this template. For users familiar with the BIG-IP, there is a manual configuration table at the end of this guide. However, because of the complexity of this configuration, we recommend using the iApp template.
- The overwhelming majority of the configuration guidance in this document is performed on F5 devices. We provide a summary of Exchange configuration steps for reference only; for complete information on how to deploy or configure the components of Microsoft Exchange Server 2010, consult the appropriate Microsoft documentation. F5 cannot provide support for Microsoft products.
- For this deployment guide, the BIG-IP LTM system **must** be running version 11.0 or later. If you are using a previous version of the BIG-IP LTM system, see the Deployment Guide index on F5.com. The configuration in this guide does not apply to previous versions.
- If you are using the BIG-IP system to offload SSL, we assume you have already obtained an SSL certificate and key, and it is installed on the BIG-IP LTM system. To configure your Client Access servers to support SSL offloading, you must first follow the Microsoft documentation. See <http://social.technet.microsoft.com/wiki/contents/articles/how-to-configure-ssl-offloading-in-exchange-2010.aspx>. Make sure you follow the correct steps for the version of Exchange Server that you are using (Exchange Server 2010 or Exchange Server 2010 SP1). If you are using Exchange Server 2010 SP2, follow the instructions for SP1.

**Important**



- The current version of the iApp template for Microsoft Exchange 2010 does not always function properly or may be slow to respond when using Internet Explorer version 7.0 or 8.0. If using the configuration described in this guide, you should use a more recent version of Internet Explorer, or an alternate web browser.
- If deploying BIG-IP Access Policy Manager (APM) features, including Edge Gateway, you must fully license and provision APM before starting the iApp template.
- If you are configuring the iApp to use BIG-IP Edge Gateway or APM, you must configure your BIG-IP system to use a DNS server that is able to resolve names in your Active Directory domain(s). You must also configure an NTP time source for your BIG-IP system

and your domain controller(s) so their times are closely synchronized. See *Appendix A: Configuring DNS and NTP settings on the BIG-IP on page 44* for instructions.

- We generally recommend that you do not re-encrypt traffic between your BIG-IP Edge Gateway and BIG-IP LTM because both BIG-IP systems must process the SSL transactions. However, if you do choose to re-encrypt, we strongly recommend you use a valid certificate (usually SAN-enabled) rather than the default, self-signed certificate for the Client SSL profile on your BIG-IP LTM system. If not re-encrypting traffic, you do not need a certificate on your BIG-IP LTM.
- This template currently only supports the use of a single DNS name and corresponding certificate and key for all services, or multiple DNS names using a SAN-enabled certificate and key. Support for multiple names, each with separate corresponding certificates and keys, will be in a future release.
- If using a single virtual server for all HTTP-based Client Access services as recommended, you **must** obtain the Subject Alternative Name (SAN) certificate and key from a 3rd party certificate authority that supports SAN certificates, and then import it onto the BIG-IP. In versions prior to 11.1, the BIG-IP does not display SAN values in the web-based Configuration utility, but uses these certificates correctly.

**Note**



For more information on SAN certificates, see *Subject Alternative Name (SAN) SSL Certificates on page 70*.

- If using BIG-IP Edge Gateway or APM, these are the *Exchange Server (Client Access Server) settings*:

Role	Out-of-the-box setting	Your Setting	Notes
SSL Offload for all HTTP services <sup>1</sup>	Not enabled	<b>Enabled</b>	Optional but strongly recommended
Client Access Array	Not configured	<b>Enabled</b>	Required
OWA Authentication <sup>1</sup>	Forms <sup>2</sup>	<b>Forms (default)</b> <sup>2</sup>	Required
Autodiscover Authentication <sup>1</sup>	Negotiate	<b>Negotiate (default)</b>	Required
ActiveSync Authentication <sup>1</sup>	Basic	<b>Basic (default)</b>	Required
Outlook Anywhere Authentication <sup>1,3</sup>	Basic	<b>Basic (default)</b>	Required

<sup>1</sup> See the following link for more information on default authentication methods for Exchange Server 2010: <http://technet.microsoft.com/en-us/library/bb331973.aspx>

<sup>2</sup> You must change the default Forms logon format from Domain\username to just username. More information is available in the OWA configuration section of this guide.

<sup>3</sup> Outlook Anywhere is disabled by default in Exchange 2010; you must enable it before you can use it.

In our example, we use the following conventions. In your configuration, you may have the same FQDN for Outlook Anywhere, OWA, and RPC Client Access, and/or use split DNS to direct internal and external clients to different virtual servers:

outlook.example.com	FQDN for Outlook Anywhere
owa.example.com	FQDN for all other HTTP services
mapi.example.com	FQDN for Client Access Array
192.0.2.0/24	Network configured for external use on the BIG-IP EDGE Gateway
10.0.0.0/24	Network configured for use on the BIG-IP LTM

**Note**



Your network topology may differ considerably from the example shown. You may choose to use separate names for all four HTTP services and the RPC Client Access service (Client Access Array).

## DNS Settings

This table contains information on DNS settings (and our example settings) for this deployment.

Record	External DNS	Internal DNS
A Records	<p>owa.example.com: 192.0.2.10 outlook.example.com: 192.0.2.11</p> <p>If the SRV record listed below is not used, you must also have at least one of these, set to the same IP as your OWA FQDN:</p> <p>example.com: 192.0.2.10 autodiscover.example.com: 192.0.2.10</p>	<p>owa.example.com: 192.0.2.10 mapi.example.com: 10.0.0.10</p> <p>If the SRV record listed below is not used <b>and</b> you don't want to use the SCP, you must also have at least one of these, set to the same IP as your OWA FQDN:</p> <p>example.com: 192.0.2.10 autodiscover.example.com: 192.0.2.10</p> <p>To prevent internal users from receiving a password prompt, internal DNS must not have an A record for the FQDN for Outlook Anywhere.</p>
SRV Records	<p>_autodiscover._tcp.example.com: port 443, host 'owa.example.com'</p>	<p>_autodiscover._tcp.example.com: port 443, host 'owa.example.com' (optional; Outlook can use SCP instead. See note above and Further Reading below)</p>

Further reading:

- Summary of SRV records on Wikipedia: [http://en.wikipedia.org/wiki/SRV\\_record](http://en.wikipedia.org/wiki/SRV_record)
- Specification for SRV records (RFC2782): <http://tools.ietf.org/html/rfc2782>
- Microsoft KB article on SRV records and the Autodiscover service: <http://support.microsoft.com/kb/940881>
- 'Understanding the Autodiscover Service' (including SCP information): <http://technet.microsoft.com/en-us/library/bb124251.aspx>

To provide feedback on this deployment guide or other F5 solution documents, contact us at [solutionsfeedback@f5.com](mailto:solutionsfeedback@f5.com).

## Products and versions tested

Product	Version
BIG-IP system	11.0, 11.0.1, 11.1 (downloadable iApp only)
Microsoft Exchange Server	2010, 2010 SP1, 2010 SP2


## Deployment Scenarios

The iApp greatly simplifies configuring the BIG-IP system for Microsoft Exchange 2010 Client Access Server roles. Before beginning the Application template, you must make a decision about the scenario in which you are using BIG-IP system for this deployment. The iApp presents the following three deployment options. You will choose one of these options when you begin configuring the iApp.

### This BIG-IP LTM will load balance and optimize Client Access Server traffic

You can select this scenario to manage, secure, and optimize client-generated Client Access Server traffic using the BIG-IP system. This is the traditional role of the BIG-IP LTM and should be used in scenarios where you are not deploying Edge Gateway or Access Policy Manager (APM) on a separate BIG-IP system. In this scenario, you have the option of configuring the BIG-IP APM to secure HTTP-based virtual servers on this system.

You would not select this option if you intend to deploy a separate BIG-IP Edge Gateway or APM that will provide secure remote access to Exchange CAS HTTP services.




1. All Exchange 2010 Client Access traffic goes to the BIG-IP LTM.
2. You can optionally use the BIG-IP APM module to provide secure access and proxied authentication for HTTP-based Client Access services: Outlook Web App, Outlook Anywhere, ActiveSync, and Autodiscover).
3. The BIG-IP LTM load balances and optimizes the traffic to the Client Access Servers, including the services which are not HTTP-based: RPC Client Access (MAPI), POP3, and IMAP4.

### This BIG-IP LTM will receive HTTP-based Client Access traffic forwarded by a BIG-IP Edge Gateway or APM

You can select this scenario to configure BIG-IP LTM with a single virtual server that receives Exchange Client Access HTTP-based traffic that has been forwarded by a separate BIG-IP Edge Gateway or APM. The virtual server can also accommodate direct traffic, e.g. internal clients that do not use the BIG-IP Edge Gateway, and non-HTTP traffic that is not handled by BIG-IP Edge Gateway such as POP3 and IMAP4.

This scenario would be used together with the following scenario, in which you configure a separate BIG-IP APM or Edge Gateway to send traffic to this BIG-IP LTM device.



1. The BIG-IP LTM receives HTTP-based Client Access traffic from a separate Edge Gateway or APM, or directly received the non HTTP-based traffic.

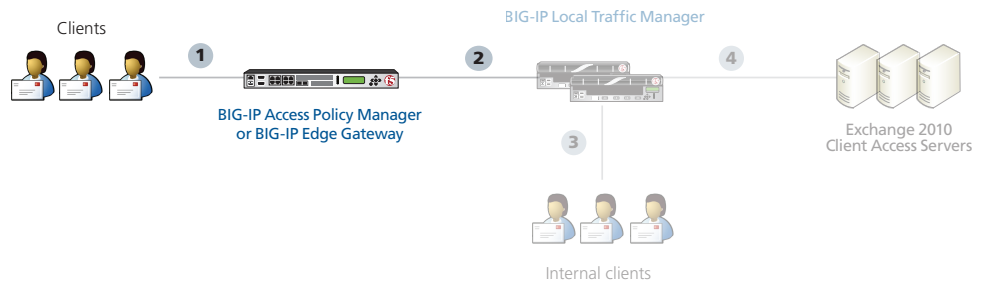
2. If you have internal Exchange clients, all Client Access Server traffic from the internal clients goes directly to the BIG-IP LTM.

☞ *While this scenario can accommodate internal clients, we do not recommend using this virtual server in that way. We strongly recommend creating a second instance of the iApp on this BIG-IP LTM for the direct traffic/internal users. You must use a unique virtual server IP address; all of the other settings can be identical. Once both iApps have been created, you would configure Split DNS (use the same domain name, but different zones and IP addresses for internal and external clients). For more information about Split DNS, refer to your DNS documentation.*

### This BIG-IP Edge Gateway or APM will provide secure remote access to CAS

You can select this scenario to configure the BIG-IP as a BIG-IP Edge Gateway or APM that will use a single virtual server to provide proxy authentication and secure remote access to Exchange HTTP-based Client Access services without requiring the use of an F5 Edge Client. The traffic will be forwarded to another BIG-IP running LTM which provides advanced load balancing, persistence, monitoring and optimizations for those services.

This scenario would be used together with the previous scenario, in which you configure a separate BIG-IP LTM to receive traffic from this Edge Gateway or APM device.



1. HTTP-based Client Access traffic goes to the BIG-IP APM or Edge Gateway, which provides proxy authentication and secure remote access.

Note: If you want to allow RPC Client Access, POP3 or IMAP4 access from external users, you must separately configure your BIG-IP by re-running the iApp, selecting the first scenario ("This BIG-IP LTM will load balance and optimize Client Access Server traffic"), choosing which of those protocols you wish to allow, and then configuring your Client Access servers as pool members.

2. After authentication, the BIG-IP APM or Edge Gateway sends the traffic to a separate BIG-IP LTM for intelligent traffic management.

Guidance specific to each deployment scenario is contained later in this document.

## Preparation worksheets

For each section of the iApp Template, you need to gather some information, such as Client Access server IP addresses and domain information. The worksheets do not contain every question in the template, but rather include the information that is helpful to have in advance. Use the worksheet(s) applicable to your configuration. More information on specific template questions can be found on the individual pages. You might find it useful to print these tables and then enter the information.

BIG-IP LTM Preparation worksheet		
<b>Traffic arriving to this BIG-IP is:</b>	<b>Encrypted</b>	<b>Unencrypted</b>
	SSL Certificate: Key: If re-encrypting traffic to the Client Access Servers and not using the BIG-IP default certificate and key: Certificate: Key:	If encrypting traffic to the Client Access Servers and not using the BIG-IP default certificate and key: Certificate: Key:
<b>BIG-IP virtual servers and Client Access Servers will be on:</b>	<b>Same Subnet</b>	<b>Different Subnets</b>
	If the maximum number of expected concurrent users per Client Access Server is more than 6,000, you need one IP address for each 6,000 users or fraction thereof:	If the Client Access Servers are a different subnet from the BIG-IP virtual servers, and do <b>not</b> use the BIG-IP as their default gateway:  If the maximum number of expected concurrent users per Client Access Server is more than 6,000, you need one IP address for each 6,000 users or fraction thereof:
<b>Single virtual IP address for all Client Access Services or multiple addresses</b>	<b>Single virtual IP address</b>	<b>Different virtual IP addresses for different services</b>
	IP address for the BIG-IP virtual server:	You need a unique IP address for each of the Client Access services you are deploying: Outlook Web App: Outlook Anywhere: ActiveSync: Autodiscover: RPC Client Access (MAPI): POP3: IMAP4:
<b>All Client Access services handled by the same set of servers, or different Servers for different services</b>	<b>Same set of Client Access Servers for all services</b>	<b>Different Client Access Servers for different services</b>
	IP addresses of the Client Access Servers:	IP addresses for Client Access Servers for each service deploying:  Outlook Web App:                      Outlook Anywhere:  ActiveSync:                                      Autodiscover:  RPC Client Access (MAPI):              POP3:  IMAP4:
<b>Outlook Web App URI</b>	If you are deploying Outlook Web App, what is the URI for reaching OWA if different than the default (http(s)://<fqdn>/owa/):	
<b>RPC Client Access ports</b>	If you are deploying RPC Client Access (MAPI), and do not want to use the default Dynamic port range, specify a port for: <b>MAPI:</b> <b>Address Book:</b>	

BIG-IP LTM Preparation Worksheet (continued): Server health monitor configuration

<p><b>Advanced Monitor configuration</b></p>	<p>If you want the iApp to configure advanced health monitors which perform logins to HTTP-based, POP3, and IMAP4 Client Access services (as opposed to simple monitors which only check network connectivity), you need the following information:</p> <p>If deploying Autodiscover, email address for monitoring:</p> <p>Mailbox account name in Active Directory for the monitors:</p> <p>Associated password:</p> <p>Domain name (can be FQDN or NETBIOS) of the user account used for monitors:</p>	<p>Second mailbox for monitoring (recommended): If deploying Autodiscover, 2<sup>nd</sup> email address for monitoring:</p> <p>2<sup>nd</sup> mailbox account name in Active Directory for the monitors:</p> <p>Associated password for this account:</p> <p>2<sup>nd</sup> domain name (can be FQDN or NETBIOS) of the user account used for monitors:</p>
<p><b>Outlook Web App authentication method</b></p>	<p>If deploying Outlook Web App, which authentication method have you configured: Forms-Based Authentication (default) <b>Important Note:</b> If you are deploying APM or Edge Gateway, you <b>must</b> use Forms-Based. Basic or Windows Integrated authentication</p>	
<p><b>Same FQDN for all HTTP-based Client Access services or different FQDNs</b></p>	<p style="text-align: center;"><b>Same FQDN</b></p> <p>FQDN for all HTTP-based Client Access services:</p>	<p style="text-align: center;"><b>Different FQDNs</b></p> <p>You need a FQDN for each HTTP-based Client Access services you are deploying:</p> <p>Outlook Web App: Outlook Anywhere: ActiveSync: Autodiscover:</p>

BIG-IP Access Policy Manager Preparation Worksheet

<p><b>Outlook Web App FQDN</b></p>	<p>If you are deploying BIG-IP APM and Outlook Web App, you need the FQDN this is used to access OWA (such as owa.example.com):</p>
<p><b>Active Directory name or IP address that BIG-IP can contact</b></p>	<p>What is the Active Directory name or IP address that this BIG-IP can contact (if name, use FQDN and not NETBIOS name):</p>
<p><b>Active Directory Domain name for Exchange users</b></p>	<p>What is the Active Directory Domain name (must be in FQDN format):</p>
<p><b>Active Directory Anonymous binding</b></p>	<p>If Anonymous Binding is not allowed in your Active Directory implementation, you need an Active Directory account with administrative permissions: User name: Password:</p>

BIG-IP Edge Gateway Preparation Worksheet

<p><b>Edge Gateway virtual server</b></p>	<p>What is the IP address you want to use for your BIG-IP Edge Gateway virtual server:</p>
<p><b>SSL Certificate and Key</b></p>	<p>SSL Certificate: Key:</p>
<p><b>Re-encrypt the traffic to the BIG-IP virtual server</b></p>	<p>You must know if the remote BIG-IP LTM that will receive traffic from this Edge Gateway is using the a self-signed/default certificate and key or a certificate signed by a Certificate Authority.</p>
<p><b>Remote LTM virtual server</b></p>	<p>What is the virtual server address on the remote BIG-IP LTM to which this Edge Gateway will forward traffic:</p>
<p><b>Outlook Web App URI</b></p>	<p>If you are deploying Outlook Web App, what is the URI for reaching OWA if different than the default (http(s)://&lt;fqdn&gt;/owa/). If you have configured OWA to use the root directory (such as http(s)://&lt;fqdn&gt;/ rather than the default http(s)://&lt;fqdn&gt;/owa/), you should <b>not</b> change this value:</p>

## Downloading and importing the new iApp template

The first task is to download and import the new Exchange Server 2010 Client Access Server iApp template. Future versions of the BIG-IP system will contain this iApp by default.

### To download and import the iApp

1. Open a web browser and go to <http://support.f5.com/kb/en-us/solutions/public/13000/400/sol13497.html>.
2. Follow the instructions to download the Microsoft Exchange 2010 iApp from downloads.f5.com to a location accessible from your BIG-IP system.  
*You must download the file, and not copy and paste the contents. F5 has discovered the copy paste operation does not work reliably.*
3. Extract (unzip) the **microsoft\_exchange\_2010\_cas.2012\_04\_06.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.

**Important**



## Configuring the BIG-IP iApp for Microsoft Exchange Server 2010

Use the following guidance to help you configure the BIG-IP system for Microsoft Exchange Server 2010 Client Access Role using the BIG-IP iApp template.

### Getting Started with the iApp for Exchange 2010

To begin the Exchange 2010 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 **Exchange-2010\_**.
5. From the **Template** list, select **f5.microsoft\_exchange\_2010\_cas.2012\_04\_06**. The new Microsoft Exchange 2010 template opens.

## Advanced options

If you select **Advanced** from the **Template Selection** list, 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.

## Analytics

This section of the template asks questions about Analytics. The Application Visibility Reporting (AVR) module allows you to view statistics specific to your Microsoft Exchange 2010 implementation. AVR is available on all BIG-IP systems v11 and later, however you must have the AVR provisioned for this option to appear. Note that these are only for application visibility reporting, you can view object-level statistics from the BIG-IP 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 template.

1. **Enable Analytics**

Choose whether you want to enable AVR for Analytics.

2. **Analytics Profile**

If you decide to use AVR, you must decide whether to use the default Analytics profile, or create a new one. As mentioned previously, we recommend creating a new profile to get the most flexibility and functionality out of AVR. If you choose to create a new profile after starting the template, you must exit the iApp, create the profile, and then restart the iApp.

To use the default Analytics profile, choose Use **Default Profile** from the list.

3. **Custom Analytics Profile**

To choose a custom profile, leave the list set to **Select a Custom Profile**, and then from the Analytics profile list, select the custom profile you created.

## Which scenario describes how you will use the BIG-IP in your CAS deployment?

Choose the option that best describes how you plan to use the BIG-IP system you are currently configuring. The scenario you select from the list determines the questions that appear in the rest of the iApp. These scenarios were described in *Deployment Scenarios on page 5*.

**Note** →

Guidance for each scenario is contained in a separate section of this deployment guide. Click the link to go to the relevant section of the guide for the scenario you plan to deploy.

1. **LTM will load balance and optimize CAS traffic**

Select this scenario to manage, secure, and optimize client-generated Client Access Server traffic using the BIG-IP system. This is the traditional role of the BIG-IP LTM and should be used in scenarios where you are not deploying Edge Gateway or APM on a separate BIG-IP system.

In this scenario, if you have fully licensed and provisioned BIG-IP APM you have the option of configuring it to provide proxy authentication for HTTP-based services on this system.

Do not select this option if you intend to deploy a separate BIG-IP Edge Gateway or APM that will provide secure remote access to Exchange CAS HTTP-based services.

If you choose this role, go to [Configuring the BIG-IP LTM to load balance and optimize Client Access Server traffic on page 12](#).

2. **LTM will receive HTTP-based CAS traffic forwarded by a BIG-IP Edge Gateway or APM**

Select this scenario to configure BIG-IP LTM with a single virtual server that receives Exchange Client Access HTTP-based traffic that has been forwarded by an BIG-IP Edge Gateway or APM. The virtual server can also accommodate direct traffic, for example internal clients that do not use the BIG-IP Edge Gateway, and non-HTTP traffic that is not handled by BIG-IP Edge Gateway such as POP3 and IMAP4.

If you choose this role, go to [Configuring the LTM to receive HTTP-based Client Access traffic forwarded by a BIG-IP Edge Gateway or APM on page 23](#).

3. **BIG-IP Edge Gateway or APM will provide secure remote access to CAS**

Select this role to configure the BIG-IP as an BIG-IP Edge Gateway or APM that will use a single HTTPS (port 443) virtual server to provide proxy authentication and secure remote access to Exchange HTTP-based Client Access services without requiring the use of an F5 Edge Client. The traffic will be forwarded to another BIG-IP running LTM which will provide advanced load balancing, persistence, monitoring and optimizations for those services.

If you choose this role, go to [Configuring the BIG-IP Edge Gateway or APM to provide secure remote access to Client Access Servers on page 32](#).

## Configuring the BIG-IP LTM to load balance and optimize Client Access Server traffic

If you chose the first scenario, *LTM will load balance and optimize CAS traffic*, use this section for guidance on configuring the iApp. Again, you should not use this option if you intend to deploy a separate BIG-IP Edge Gateway or APM that will provide secure remote access to the HTTP-based Client Access services.

### Deploying APM?

The first section of the iApp in this scenario asks about the BIG-IP Access Policy Manager. You must have APM fully licensed and provisioned to use APM. If you are not deploying APM, continue with the next section. For more information on BIG-IP APM, see <http://www.f5.com/products/big-ip/access-policy-manager.html>.

As mentioned in the prerequisites, if you are deploying APM, you must have configured the BIG-IP system for DNS and NTP. See *Appendix A: Configuring DNS and NTP settings on the BIG-IP on page 44* for instructions.

1. **Do you want to deploy Access Policy Manager?**

If you want to use the BIG-IP Access Policy Manager to provide proxy authentication and secure remote access for HTTP-based Client Access services, select **Yes** from the list. If you select Yes, additional questions appear.

If you do not want to deploy APM at this time, select **No** and continue with the next section. If you decide add APM at a later time, you can always re-enter the iApp and then select Yes.

2. **Outlook Web App FQDN**

If you are deploying Outlook Web App (OWA), type the fully qualified domain name (FQDN) that clients use to access OWA. This is the FQDN you must configure in DNS to resolve to the IP address of the BIG-IP virtual server providing services for Outlook Web App.

Do not include the protocol (such as "https://") or any path information (such as "/owa"). If you are not deploying OWA at this time, you may leave this field blank.

3. **Active Directory name or IP address**

Type the name or IP address of an Active Directory server in your domain that this BIG-IP system can contact. If you provide a name rather than an IP address, use the FQDN of the server rather than the NetBIOS name. Make sure this BIG-IP system and the Active Directory server have routes to one another and that firewalls allow traffic between the two.

4. **Active Directory domain name**

Type the Active Directory domain name in FQDN format.

5. **Active Directory Anonymous binding**

If anonymous binding is allowed, no additional information is required.

If your Active Directory implementation does not allow anonymous binding, select **Credentials are required for binding**. Two new questions appear.

A. *Active Directory user name*

Type a user name with administrative permissions.

B. *Password*

Type the associated password.

**Note**

*These credentials are stored in plain text on your BIG-IP system. This completes the BIG-IP APM section.*

## Tell us about your deployment

In this section, the iApp gathers general information about your Client Access Server deployment. Remember, you must import an SSL certificate and key that correspond to all fully-qualified DNS names that you are using for OWA, Outlook Anywhere, Autodiscover, ActiveSync, POP3, or IMAP4 traffic. Any certificate that you obtain with multiple names must be in SAN (Subject Alternative Name) format, not SNI (Server Name Indication) format.

### 1. ***Incoming traffic encrypted or unencrypted***

Select whether any of the HTTP-based, POP3 and IMAP4 traffic will be encrypted or not when it arrives on this system. In nearly all cases for this deployment scenario, it will be encrypted (it would not be encrypted, for example, if you selected one of the other scenarios/roles for this iApp, and elected to offload SSL/TLS traffic at a separate BIG-IP Edge Gateway or APM).

Note that the BIG-IP does not offload the encryption used for RPC; the answer to this question should be based on the other Client Access protocols you intend to deploy.

**Note**



*This question does not appear if you chose to deploy APM in the previous section. If you selected to deploy APM, continue with the SSL Certificate questions (A:i, ii, iii) below.*

#### A. **Encrypted**

If you chose Encrypted in the previous question, additional questions appear.

##### i. SSL Certificate

Select the SSL certificate you imported onto the BIG-IP system for decrypting client connections.

If you have not yet imported a certificate, you can leave the default selections and reconfigure this iApp after obtaining the certificates. The deployment will not function correctly until you have selected the correct certificates here.

**Note**



*For detailed information on SAN certificates, see [Subject Alternative Name \(SAN\) SSL Certificates on page 70](#).*

##### ii. SSL Key

Select the associated key from the list.

##### iii. Re-encrypt?

If you want the BIG-IP system to offload SSL processing from the Client Access Servers, select **Do not re-encrypt (SSL Offload)** from the list. Offloading SSL on the BIG-IP system can extend Exchange Server 2010 server capacity.

If you choose SSL Offload, you must have followed the instructions in <http://social.technet.microsoft.com/wiki/contents/articles/how-to-configure-ssl-offloading-in-exchange-2010.aspx> as described in the prerequisites.

If you choose **Re-encrypt (SSL Bridging)**, the BIG-IP system encrypts the traffic headed for the Client Access Servers.

#### B. **Unencrypted**

If Client Access traffic is arriving at this BIG-IP system unencrypted (typically because you configured to offload SSL/TLS traffic at the BIG-IP Edge Gateway or APM that is sending Client Access traffic to this device) select **Unencrypted** from the list.

i. Encrypt the traffic?

If you chose Unencrypted in the previous question, this question asks if you want to encrypt the traffic to the Client Access Servers.

If you choose **Do not encrypt (SSL Offload)**, the BIG-IP system does not modify the traffic, and you can continue with the next question.

If you choose **Encrypt (SSL Bridging)**, the BIG-IP system encrypts the traffic headed for the Client Access Servers.

2. **WAN or LAN**

Choose whether most clients are connecting over a WAN or LAN. The iApp uses your selection to configure the proper TCP optimization settings.

3. **Location of BIG-IP virtual servers in relation to Client Access Servers**

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

A. **Same subnet for BIG-IP virtual servers and Client Access Servers**

If the BIG-IP virtual servers and Client Access Servers are on the same subnet, SNAT is configured on the BIG-IP virtual server and you must specify the number of concurrent users in a question below.

B. **Different subnet for BIG-IP virtual servers and Client Access Servers**

If the BIG-IP virtual servers and Client Access Servers are on different subnets, the following question appears asking how routing is configured.

i. Routing configuration

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

If the Client Access Servers do not use the BIG-IP as their default gateway, SNAT is configured on the BIG-IP virtual server and you must select the expected number of concurrent users in the next question.

If the Client Access Servers use the BIG-IP as their default gateway, the concurrent user question does not appear.

4. **Maximum number of concurrent users**

Select whether you expect more than 6,000 concurrent users for each Client Access Server. This is used to determine whether the BIG-IP system uses a SNAT Pool or SNAT Automap for your Client Access Server deployment.

**Note**



For specific information on SNAT Pools, including why we chose 6,000 concurrent users per Client Access Server, see *Maximum number of concurrent users: SNAT Pool guidance on page 70*.

A. **Less than 6000 per Client Access Server**

If you expect fewer than 6,000 concurrent users, choose **Less than 6000**. The iApp uses SNAT Automap, and no additional addresses are required.

B. **More than 6000 per Client Access Server**

If you expect more than 6,000 concurrent users, choose **More than 6000**. The iApp creates a SNAT pool and a new question appears asking for at least one available IP address for every 6,000 users, or fraction thereof.

In the **IP** box, type an IP address. The IP address(es) you specify must **not** be self IP addresses of this BIG-IP system.

Click **Add** to include additional addresses.

**Important**

*If you choose more than 6,000 users, and do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per Client Access server is reached, new requests will fail.*

This question does not appear if you chose different subnets for BIG-IP virtual servers and Client Access Servers **and** the Client Access Servers use the BIG-IP as their default gateway.

5. **Single IP address or multiple IP addresses for Client Access Server connections**

Choose whether you want to use a single IP address for all Client Access connections, or separate IP addresses for the different services. If you chose a single IP address, the iApp creates a single virtual server for all of the Client Access services. If you choose different addresses, the BIG-IP creates individual virtual servers for each service. There are advantages to each method:

A. **Single IP address**

With a IP address, you can combine multiple functions on the same virtual server; for instance, you may wish to have a single fully-qualified domain name (FQDN) and associated SSL certificate for all HTTP-based Client Access methods. You only need to provision a single IP address for the virtual server. If you want the services to have unique DNS names despite sharing an IP address, you need to obtain an SSL certificate that supports Subject Alternative Names. For detailed information on SAN certificates, see [\*Subject Alternative Name \(SAN\) SSL Certificates on page 70\*](#).

B. **Different IP addresses for different services**

By maintaining a separate virtual server for each component, you can manage each service largely independently from one another. For instance, you may wish to have different pool membership, load balancing methods, or custom monitors for Outlook Web App and Outlook Anywhere. If each of those services are associated with a different virtual server, granular management becomes easier. You need to provision an available IP address for each virtual server, and obtain a valid SSL certificate with a unique subject name for each service.

6. **Distribution of Client Access protocols**

Choose whether all your Client Access services are handled by the same Client Access Servers, or if each service is handled by a unique set of Client Access Servers.

This iApp creates separate pools and monitors for each service regardless of this setting. However, if you use the same set of servers for all services, you only have to specify the server IP addresses once.

## Tell us about which services you are deploying

In this section, the iApp gathers information about which Client Access services you are deploying.

1. ***Customize BIG-IP pool settings***

Choose whether you want to customize the BIG-IP load balancing pools for the Client Access services, or use the F5 recommended settings.

If you select **Use settings recommended by F5**, continue with #2.

If you select **Customize Pool Settings**, the following options appear:

A. **Which load balancing method do you want to use**

Choose the load balancing method you want to use. We recommend the default, Least Connections (member). See the BIG-IP documentation for a description of each load balancing method.

B. **Do you want the BIG-IP 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 Ask F5.

**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 Client Access Server nodes.*

If you want the BIG-IP to queue TCP requests, select **Yes** from the list. Additional questions appear.

i. *Maximum number of TCP requests to be queued*

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.

ii. *How long before requests expire*

Type a number of milliseconds for the TCP request timeout value.

2. ***BIG-IP virtual server IP address***

This question appears if you chose a single IP address for Client Access Server connections in question #5 on *page 15*. If you chose a single IP address, type a valid IP address to use for the BIG-IP virtual servers.

3. ***Deploying OWA (including ECP)***

Select whether you are deploying Outlook Web Access at this time. This includes the Exchange Control Panel (ECP). If you are deploying these services at this time, select **Yes**.

If you select Yes, additional questions may appear (depending on your previous selections).

A. **URI for reaching Outlook Web App**

If you are deploying Outlook Web App, you must specify the URI for reaching OWA. By default, the URI is **/owa/**. If your URI is different, type it here.

**Note**

*We do not recommend changing the default URI for Outlook Web App. If you modified the URI to use the root directory (**http(s)://<fqdn>/**), you must NOT change this value from the default.*

B. **IP address for the Outlook Web App virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the Outlook Web App virtual server. Type the IP address here.

If you selected a single IP address for Client Access services, this question does not appear.

C. **IP addresses of the Outlook Web App servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the Outlook Web App servers. Type the IP addresses. Click the Add button to add additional servers.

If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

4. **Deploying Outlook Anywhere (includes EWS and OAB)**

Select whether you are deploying Outlook Anywhere at this time. This includes EWS (Exchange Web Services) and OAB (Offline Address Book).

If you select Yes, additional questions may appear depending on your previous selections.

**Important**

*You must enable Outlook Anywhere on each of your Exchange Client Access Servers before that service will be available. Outlook Anywhere is not enabled by default on Exchange Client Access Servers. See the Microsoft documentation for specific instructions.*

*To prevent internal users from receiving a password prompt, your internal DNS must not have an 'A' record for the FQDN for Outlook Anywhere.*

A. **IP address for the Outlook Anywhere virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the Outlook Anywhere virtual server. Type the IP address here.

If you selected a single IP address for Client Access services, this question does not appear.

B. **IP addresses of the Outlook Anywhere servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the Outlook Anywhere servers. Type the IP addresses. Click the Add button to add additional servers.

If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

5. ***Deploying ActiveSync***

Select whether you are deploying ActiveSync at this time.

If you select Yes, additional questions may appear (depending on your previous selections).

A. **IP address of the ActiveSync virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the ActiveSync virtual server. Type the IP address here.

If you selected a single IP address for Client Access services, this question does not appear.

B. **IP addresses for the ActiveSync servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the ActiveSync servers. Type the IP addresses. Click the Add button to add additional servers.

If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

6. ***Deploying Autodiscover***

Select whether you are deploying Autodiscover at this time.

If you select Yes, additional questions may appear (depending on your previous selections).

**Critical**



*To deploy Autodiscover, you must either create an 'SRV' record in DNS or create 'A' records in order for external clients to be able to make use of Autodiscover. If you do not want to use an 'SRV' record, then you must have 'A' records for either 'autodiscover.<yourdomain>' or '<yourdomain>' that resolve to the IP address you have designated for your Autodiscover virtual server.*

A. **IP address for the Autodiscover virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the Autodiscover virtual server. Type the IP address here.

If you selected a single IP address for Client Access services, this question does not appear.

B. **IP addresses of the Autodiscover servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the Autodiscover servers. Type the IP addresses. Click the Add button to add additional servers.

If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

7. ***Deploying RPC Client Access (MAPI)***

Select whether you are deploying RPC Client Access (MAPI) at this time.

If you select Yes, additional questions appear.

**Critical**



*You must enable and configure a Client Access Array in your Exchange Server site before RPC Client Access will function. See [Creating a new Client Access Array on page 72](#) for more information.*

*If deploying RPC Client Access, you must also deploy Outlook Anywhere, to properly handle EWS (Exchange Web Services) traffic.*

A. **Static ports or dynamic port range**

If you choose the default dynamic range of ports, no additional information is necessary, continue with the following question.

If you chose Static Ports, additional questions appear:

i. MAPI Port

Type the port you are using for MAPI.

ii. Address Book

Type the port you are using for the Address book.

**Important**

*You must make sure each of your Client Access Servers are configured to use the static ports you specified here. See <http://social.technet.microsoft.com/wiki/contents/articles/configure-static-rpc-ports-on-an-exchange-2010-client-access-server.aspx> for more information.*

B. **IP address for the RPC Client Access virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the RPC Client Access virtual server. Type the IP address here. If you selected a single IP address for Client Access services, this question does not appear.

C. **IP addresses of the RPC Client Access servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the RPC Client Access servers. Type the IP addresses. Click the Add button to add additional servers. If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

8. **Deploying POP3**

Select whether you are deploying POP3 at this time.

If you select Yes, additional questions may appear (depending on your previous selections).

**Important**

*You must enable POP3 on each of your Exchange Client Access Servers before that service will be available. POP3 is not enabled by default on Exchange Client Access Servers.*

*If you are offloading SSL, you must configure the Authentication properties for POP3 on each of your Exchange Client Access Servers to allow logins using plain text. By default, POP3 is configured to only allow secure (encrypted) logins.*

A. **IP address for the POP3 virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the POP3 virtual server. Type the IP address here. If you selected a single IP address for Client Access services, this question does not appear.

B. **IP addresses of the POP3 servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the POP3 servers. Type the IP addresses. Click the Add button to add additional servers. If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit. If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

**Important**



9. ***Deploying IMAP4***

Select whether you are deploying IMAP4 at this time.

If you select Yes, additional questions may appear (depending on your previous selections).

*You must enable IMAP4 on each of your Exchange Client Access Servers before that service will be available. IMAP4 is not enabled by default on Exchange Client Access Servers.*

*If you are offloading SSL, you must configure the Authentication properties for IMAP4 on each of your Exchange Client Access Servers to allow logins using plain text. By default, IMAP4 is configured to only allow secure (encrypted) logins.*

A. **IP address for the IMAP4 virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the IMAP4 virtual server. Type the IP address here.

If you selected a single IP address for Client Access services, this question does not appear.

B. **IP addresses of the IMAP4 servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the IMAP4 servers. Type the IP addresses. Click the Add button to add additional servers.

If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

10. ***IP Addresses of your Client Access Servers***

If you chose a single IP address for all Client Access services and that each Client Access service will be handled by the same Client Access Servers, the iApp asks for the IP addresses of the Client Access Servers. Type the IP addresses. Click the Add button to add additional servers.

If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

## Server Health Monitors

The last section of the template asks for information about the health checks the iApp will configure for the Client Access Servers.

1. ***Health monitor interval***

Specifies how often the system checks the health of the servers. We recommend the default of 30 seconds.

2. ***Advanced or simple health monitors***

Choose whether you want to use advanced or simple health monitors:

A. **Use simple monitors**

Simple monitors check network connectivity but do not perform actual logins. If you use simple monitors, the BIG-IP LTM may not be able to completely determine status of Client Access services.

**Important**



*If you select simple monitors and are configuring the iApp for POP3 and/or IMAP, you must add a user name and password to the monitor after completing the iApp. See [Optional: Adding a user name and password to POP3 and IMAP4 simple monitors on page 40](#).*

**B. Use advanced monitors**

If you choose advanced monitors, the BIG-IP performs logins to most of the Client Access services (all except RPC/MAPI) and checks for valid content in the response. Because these monitors attempt to access a specific mailbox, they can more accurately determine the actual health of Client Access services. However, account maintenance and Mailbox status must become a part of your monitoring strategy. For example, if an account used for monitoring is locked or deleted, the monitor will mark the services down for all users.

We strongly recommend creating a mailbox account(s) specifically for use in the monitor(s). The accounts for those mailboxes should have no other privileges in the domain and should be configured with passwords that do not expire.

If you choose advanced monitors, additional questions appear:

- i. *Email address for Autodiscover*  
If you chose to deploy Autodiscover, type the email address associated with the account you are going to monitor (and that you specify in the following question).
- ii. *Mailbox account for monitors*  
Type a mailbox account for use in the advanced monitors. This name corresponds to the account name field in Active Directory (rather than the email address).
- iii. *Password*  
Type the associated password. Note that credentials are stored in plain text on this BIG-IP system.
- iv. *Domain*  
Type the Domain name for the user account. This domain can be entered in either FQDN (mydomain.example.com) or NetBIOS (MYDOMAIN) format.
- v. *Monitor a second mailbox*  
Choose whether you want to monitor a second mailbox. We strongly recommend configuring a second mailbox account to be used by a second set of monitors, using a mailbox that is configured to reside on a different Mailbox server. BIG-IP LTM will only make a Client Access service on a specific server down if both sets of credentials fail. This provides resiliency to accommodate configuration errors with a single account, mailbox, or Mailbox server.

3. **Authentication method for Outlook Web App**

If you configured the iApp to deploy Outlook Web App at this time, choose the authentication method you have configured for Outlook Web App. You can choose either **Forms-Based** (default) or **Basic/Windows Integrated Authentication**. The health monitors will be customized to accommodate the authentication method you are using.

**Important**



*If you are using APM in this scenario, you must choose **Forms-Based**. If you are using Forms-Based authentication for OWA, you must change the credential format required for OWA on each Exchange Client Access Server from the default domain and username format to just the username.*

4. **Same FQDN for all HTTP-based Client Access Services**

Choose whether you are using one FQDN for all HTTP-based services or separate FQDNs for each service. These values are used for HTTP 1.1-based health monitors.

A. **One FQDN for all HTTP services**

If you are using one FQDN, choose **One FQDN for all HTTP services**.

In the next box, type the FQDN you are using for your HTTP-based Client Access services.

B. **Different FQDNs for each HTTP service**

If you are using different FQDNs, choose **Different FQDNs for each HTTP service**.

Additional questions appear asking for each individual FQDN. Type the appropriate FQDN in the appropriate box.

### Additional Steps

Review the information in the Additional steps section, and take appropriate action if necessary. All of the notes in Additional Steps are found in the relevant section of this deployment guide.

### Finished

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

**Critical**



*There are important modifications you must make to the configuration produced by the iApp template. See [Modifying the iApp template configuration](#) on page 35.*

## Configuring the LTM to receive HTTP-based Client Access traffic forwarded by a BIG-IP Edge Gateway or APM

If you chose the second scenario, *LTM will receive HTTP-based CAS traffic forwarded by a BIG-IP Edge Gateway or APM*, use this section for guidance on configuring the iApp. This selection configures BIG-IP LTM with a single virtual server that receives Exchange Client Access HTTP-based traffic that has been forwarded by a separate BIG-IP Edge Gateway or APM. The virtual server can also accommodate non-HTTP traffic that is not handled by BIG-IP Edge Gateway such as POP3 and IMAP4.

While this virtual server can be used for direct traffic (for example, internal clients that do not use the BIG-IP Edge Gateway), we do not recommend using this virtual server in that way. For direct traffic, we strongly recommend creating a second instance of the iApp on this BIG-IP LTM for the direct traffic/internal users. You must use a unique virtual server IP address, all of the other settings can be identical. Once both iApps have been created, you would configure Split DNS (use the same domain name, but different zones and IP addresses for internal and external clients). For more information about Split DNS, refer to your DNS documentation.

### Tell us about your deployment

In this section, the iApp gathers general information about your Client Access Server deployment.

1. ***Incoming traffic encrypted or unencrypted***

Select whether any of the HTTP-based, POP3, or IMAP4 traffic will be encrypted when it arrives on this system. Because you may have configured to offload SSL/TLS traffic at the Edge Gateway or APM that is sending Client Access traffic to this device, the traffic may be arriving unencrypted.

Note that the BIG-IP does not offload the encryption used for RPC; the answer to this question should be based on the other Client Access protocols you intend to deploy.

A. **Encrypted**

If you chose Encrypted in the previous question, additional questions appear.

i. SSL Certificate

Select the SSL certificate you imported onto the BIG-IP system for decrypting client connections. This Edge Gateway/APM deployment supports the use of a single virtual server with one IP address, so if you are using different DNS names for different services, you need to provide a Subject Alternative Name (SAN) enabled certificate and key that is valid for each name.

*For detailed information on SAN certificates, see Subject Alternative Name (SAN) SSL Certificates on page 70.*

ii. SSL Key

Select the associated key from the list.

iii. Re-encrypt?

If you chose Unencrypted in the previous question, this question asks if you want to encrypt the traffic to the Client Access servers.

If you choose **Do not re-encrypt (SSL Offload)**, the BIG-IP system does not modify the traffic, and you can continue with the next question.

If you choose **Re-encrypt (SSL Bridging)**, the BIG-IP system encrypts the traffic headed for the Client Access Servers.

B. **Unencrypted**

If Client Access traffic is arriving at this BIG-IP system unencrypted (typically because you configured to offload SSL/TLS traffic at the BIG-IP Edge Gateway or APM that is sending Client Access traffic to this device) select **Unencrypted** from the list.

i. *Encrypt the traffic?*

If you chose Unencrypted in the previous question, this question asks if you want to encrypt the traffic to the Client Access servers.

If you choose **Do not encrypt (SSL Offload)**, the BIG-IP system does not modify the traffic, and you can continue with the next question.

If you choose **Encrypt (SSL Bridging)**, the BIG-IP system encrypts the traffic headed for the Client Access Servers.

2. **WAN or LAN**

Choose whether most clients are connecting over a WAN or LAN. This setting is used to configure the proper TCP optimization settings.

3. **Location of BIG-IP virtual servers in relation to Client Access Servers**

Select whether your BIG-IP virtual servers are on the same subnet as your Client Access Servers, or on different subnets. This setting is used to determine the SNAT and routing configuration.

A. **Same subnet for BIG-IP virtual servers and Client Access Servers**

If the BIG-IP virtual servers and Client Access Servers are on the same subnet, SNAT is configured on the BIG-IP virtual server and you must specify the number of concurrent users in question #4 below.

B. **Different subnet for BIG-IP virtual servers and Client Access Servers**

If the BIG-IP virtual servers and Client Access Servers are on different subnets, the following question appears asking how routing is configured.

i. *Routing configuration*

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

If the Client Access Servers do not use the BIG-IP as their default gateway, SNAT is configured on the BIG-IP virtual server and you must select the expected number of concurrent users in the next question.

If the Client Access Servers use the BIG-IP as their default gateway, the concurrent user question does not appear.

4. **Maximum number of concurrent users**

Select if you expect more or less than 6,000 concurrent users for each Client Access Server.

A. **Less than 6000 per Client Access Server**

If you expect fewer than 6,000 concurrent users, choose **Less than 6000**. The iApp uses SNAT Automap, and no additional addresses are required.

B. **More than 6000 per Client Access Server**

If you expect more than 6,000 concurrent users, choose **More than 6000**. The iApp

creates a SNAT pool and a new question appears asking for at least one available IP address for every 6,000 users, or fraction thereof.

For specific information on SNAT Pools, including why we chose 6,000 concurrent users per Client Access Server, see *Maximum number of concurrent users: SNAT Pool guidance on page 70*

In the **IP** box, type an IP address. Click **Add** to include additional addresses.

**Important**

*If you choose more than 6,000 users, and do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per Client Access server is reached, new requests will fail.*

This question does not appear if you chose different subnets for BIG-IP virtual servers and Client Access Servers **and** the Client Access Servers use the BIG-IP as their default gateway.

5. ***Distribution of Client Access protocols***

Choose whether all your Client Access services are handled by the same Client Access Servers, or if each service is handled by a unique set of Client Access Servers.

This iApp creates separate pools and monitors for each service regardless of this setting. However, if you use the same set of servers for all services, you only have to specify the server IP addresses once.

**Tell us about which services you are deploying**

In this section, the iApp gathers information about which Client Access services you are deploying.

1. ***Customize BIG-IP pool settings***

Choose whether you want to customize the BIG-IP load balancing pools for the Client Access services, or use the F5 recommended settings.

If you select **Use settings recommended by F5**, continue with #2.

If you select **Customize Pool Settings**, the following options appear:

A. **Which load balancing method do you want to use**

Choose the load balancing method you want to use. We recommend the default, Least Connections (member). See the BIG-IP documentation for a description of each load balancing method.

B. **Do you want the BIG-IP 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 Ask F5.

**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 Client Access Server nodes.*

If you want the BIG-IP to queue TCP requests, select **Yes** from the list. Additional questions appear.

- i. Maximum number of TCP requests to be queued  
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.
- ii. How long before requests expire  
Type a number of milliseconds for the TCP request timeout value.

2. **BIG-IP virtual server IP address**

Type a valid IP address to use for the BIG-IP virtual server.

3. **Deploying OWA (includes ECP)**

Select whether you are deploying Outlook Web Access at this time. This includes the Exchange Control Panel (ECP). If you are deploying these services at this time, select **Yes**.

If you select Yes, additional questions may appear (depending on your previous selections)

A. **URI for reaching Outlook Web App**

If you are deploying Outlook Web App, you must specify the URI for reaching OWA. By default, the URI is **/owa/**. If your URI is different, type it here.

**Note**



*We do not recommend changing the default URI for Outlook Web App. If you modified the URI to use the root directory (**http(s)://<fqdn>/**), you must NOT change this value from the default.*

B. **IP address for the Outlook Web App virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the Outlook Web App virtual server. Type the IP address here.

If you selected a single IP address for Client Access services, this question does not appear.

C. **IP addresses of the Outlook Web App servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the Outlook Web App servers. Type the IP addresses. Click the Add button to add additional servers.

If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

4. **Deploying Outlook Anywhere (includes EWS and OAB)**

Select whether you are deploying Outlook Anywhere at this time. This includes EWS (Exchange Web Services) and OAB (Offline Address Book).

If you select Yes, additional questions may appear depending on your previous selections.

**Important**



You must enable Outlook Anywhere on each of your Exchange Client Access Servers before that service will be available. Outlook Anywhere is not enabled by default on Exchange Client Access Servers. See the Microsoft documentation for specific instructions.

To prevent internal users from receiving a password prompt, your internal DNS must not have an 'A' record for the FQDN for Outlook Anywhere.

**A. IP address for the Outlook Anywhere virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the Outlook Anywhere virtual server. Type the IP address here. If you selected a single IP address for Client Access services, this question does not appear.

**B. IP addresses of the Outlook Anywhere servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the Outlook Anywhere servers. Type the IP addresses. Click the Add button to add additional servers. If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

**5. Deploying ActiveSync**

Select whether you are deploying ActiveSync at this time.

If you select Yes, additional questions may appear (depending on your previous selections).

**A. IP address of the ActiveSync virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the ActiveSync virtual server. Type the IP address here. If you selected a single IP address for Client Access services, this question does not appear.

**B. IP addresses for the ActiveSync servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the ActiveSync servers. Type the IP addresses. Click the Add button to add additional servers. If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

**6. Deploying Autodiscover**

Select whether you are deploying Autodiscover at this time.

If you select Yes, additional questions may appear (depending on your previous selections).

**Critical**



To deploy Autodiscover, you must either create an 'SRV' record in DNS or create 'A' records in order for external clients to be able to make use of Autodiscover. If you do not want to use an 'SRV' record, then you must have 'A' records for either 'autodiscover.<yourdomain>' or '<yourdomain>' that resolve to the IP address you have designated for your Autodiscover virtual server.

**A. IP address for the Autodiscover virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for

the IP address to use for the Autodiscover virtual server. Type the IP address here.  
If you selected a single IP address for Client Access services, this question does not appear.

**B. IP addresses of the Autodiscover servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the Autodiscover servers. Type the IP addresses. Click the Add button to add additional servers.

If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

**7. Deploying RPC Client Access (MAPI)**

Select whether you are deploying RPC Client Access (MAPI) at this time.

If you select Yes, additional questions appear.

**Critical**



*You must enable and configure a Client Access Array in your Exchange Server site before RPC Client Access will function. See [Creating a new Client Access Array on page 72](#) for more information.*

*If deploying RPC Client Access, you must also deploy Outlook Anywhere, to properly handle EWS (Exchange Web Services) traffic.*

**A. Static ports or dynamic port range**

If you choose the default dynamic range of ports, no additional information is necessary, continue with the following question.

If you chose Static Ports, additional questions appear:

**i. MAPI Port**

Type the port you are using for MAPI.

**ii. Address Book**

Type the port you are using for the Address book.

**Important**



*You must make sure each of your Client Access Servers are configured to use the static ports you specified here. See <http://social.technet.microsoft.com/wiki/contents/articles/configure-static-rpc-ports-on-an-exchange-2010-client-access-server.aspx> for more information.*

**B. IP address for the RPC Client Access virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the RPC Client Access virtual server. Type the IP address here.

If you selected a single IP address for Client Access services, this question does not appear.

**C. IP addresses of the RPC Client Access servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the RPC Client Access servers. Type the IP addresses. Click the Add button to add additional servers.

If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

8. ***Deploying POP3***

Select whether you are deploying POP3 at this time.

If you select Yes, additional questions may appear (depending on your previous selections).

**Important**



*You must enable POP3 on each of your Exchange Client Access Servers before that service will be available. POP3 is not enabled by default on Exchange Client Access Servers.*

*If you are offloading SSL, you must configure the Authentication properties for POP3 on each of your Exchange Client Access Servers to allow logins using plain text. By default, POP3 is configured to only allow secure (encrypted) logins.*

A. **IP address for the POP3 virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the POP3 virtual server. Type the IP address here.

If you selected a single IP address for Client Access services, this question does not appear.

B. **IP addresses of the POP3 servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the POP3 servers. Type the IP addresses. Click the Add button to add additional servers.

If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

9. ***Deploying IMAP4***

Select whether you are deploying IMAP4 at this time.

If you select Yes, additional questions may appear (depending on your previous selections).

**Important**



*You must enable IMAP4 on each of your Exchange Client Access Servers before that service will be available. IMAP4 is not enabled by default on Exchange Client Access Servers.*

*If you are offloading SSL, you must configure the Authentication properties for IMAP4 on each of your Exchange Client Access Servers to allow logins using plain text. By default, IMAP4 is configured to only allow secure (encrypted) logins.*

A. **IP address for the IMAP4 virtual server**

If you selected separate IP addresses for different Client Access services, the iApp asks for the IP address to use for the IMAP4 virtual server. Type the IP address here.

If you selected a single IP address for Client Access services, this question does not appear.

B. **IP addresses of the IMAP4 servers**

If you specified that each Client Access service will be handled by a unique set of Client Access Servers, the iApp asks for the IP addresses of the IMAP4 servers. Type the IP addresses. Click the Add button to add additional servers.

If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

If you specified all services will be handled by the same set of Client Access Servers, this question does not appear.

10. ***IP Addresses of your Client Access Servers***

If you chose a single IP address for all Client Access services and that each Client Access service will be handled by the same Client Access Servers, the iApp asks for the IP addresses

of the Client Access Servers. Type the IP addresses. Click the Add button to add additional servers. If you chose to have the BIG-IP queue TCP requests, you must specify a Connection Limit.

## Server Health Monitors

The last section of the template asks for information about the health checks the iApp will configure for the Client Access Servers.

1. ***Health monitor interval***  
Specifies how often the system checks the health of the servers. We recommend the default of 30 seconds.
2. ***Advanced or simple health monitors***  
Choose whether you want to use advanced or simple health monitors:

### A. Use simple monitors

Simple monitors check network connectivity but do not perform actual logins. If you use simple monitors, the BIG-IP LTM may not be able to completely determine status of Client Access services.

### Important



If you select simple monitors and are configuring the iApp for POP3 and/or IMAP, you must add a user name and password to the monitor after completing the iApp. See *Optional: Adding a user name and password to POP3 and IMAP4 simple monitors on page 40.*

### B. Use advanced monitors

If you choose advanced monitors, the BIG-IP performs logins to most of the Client Access services (all except RPC/MAPI) and checks for valid content in the response. Because these monitors attempt to access a specific mailbox, they can more accurately determine the actual health of Client Access services. However, account maintenance and Mailbox status must become a part of your monitoring strategy. For example, if an account used for monitoring is locked or deleted, the monitor will mark the services down for all users.

We strongly recommend creating a mailbox account(s) specifically for use in the monitor(s). The accounts for those mailboxes should have no other privileges in the domain and should be configured with passwords that do not expire.

If you choose advanced monitors, additional questions appear:

- i. ***Email address for Autodiscover***  
If you chose to deploy Autodiscover, type the email address associated with the account you are going to monitor (and that you specify in the following question).
- ii. ***Mailbox account for monitors***  
Type a mailbox account for use in the advanced monitors. This name corresponds to the account name field in Active Directory (rather than the email address).
- iii. ***Password***  
Type the associated password. Note that credentials are stored in plain text on this BIG-IP system.
- iv. ***Domain***  
Type the Domain name for the user account. This domain can be entered in either FQDN (mydomain.example.com) or NetBIOS (MYDOMAIN) format.

v. *Monitor a second mailbox*

Choose whether you want to monitor a second mailbox. We strongly recommend configuring a second mailbox account to be used by a second set of monitors, using a mailbox that is configured to reside on a different Mailbox server. BIG-IP LTM will only make a Client Access service on a specific server down if both sets of credentials fail. This provides resiliency to accommodate configuration errors with a single account, mailbox, or Mailbox server.

If you choose to monitor a second mailbox, enter the Mailbox account, password, and domain information as described above.

3. ***Authentication method for Outlook Web App***

If you configured the iApp to deploy Outlook Web App at this time, choose the authentication method you have configured for Outlook Web App. You can choose either **Forms-Based** (default) or **Basic/Windows Integrated Authentication**. The health monitors will be customized to accommodate the authentication method you are using.

**Important**



*If you are using APM in this scenario, you must choose **Forms-Based**. If you are using Forms-Based authentication for OWA, you must change the credential format required for OWA on each Exchange Client Access Server from the default domain and username format to just the username.*

4. ***Same FQDN for all HTTP-based Client Access Services***

Choose whether you are using one FQDN for all HTTP-based services or separate FQDNs for each service. These values are used for HTTP 1.1-based health monitors.

A. **One FQDN for all HTTP services**

If you are using one FQDN, choose **One FQDN for all HTTP services**.

In the next box, type the FQDN you are using for your HTTP-based Client Access services.

B. **Different FQDNs for each HTTP service**

If you are using different FQDNs, choose **Different FQDNs for each HTTP service**.

Additional questions appear asking for each individual FQDN. Type the appropriate FQDN in the appropriate box.

### Additional Steps

Review the information in the Additional steps section, and take appropriate action if necessary. All of the notes in Additional Steps are found in the relevant section of this deployment guide.

### Finished

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

**Critical**



*There are important modifications you must make to the configuration produced by the iApp template. See *Modifying the iApp template configuration* on page 35.*

## Configuring the BIG-IP Edge Gateway or APM to provide secure remote access to Client Access Servers

If you chose the second scenario, *BIG-IP Edge Gateway or APM will provide secure remote access to CAS*, use this section for guidance on configuring the iApp. In this scenario, the BIG-IP will be configured as a BIG-IP Edge Gateway or APM that will use a single virtual server to provide proxy authentication and secure remote access to all Exchange HTTP-based Client Access services (Outlook Web App (including ECP), Outlook Anywhere (including EWS and OAB), ActiveSync, and Autodiscover) without requiring the use of the F5 Edge Client. The traffic will be forwarded to separate BIG-IP running LTM which will provide advanced load balancing, persistence, monitoring and optimizations for those services.

As mentioned in the prerequisites, because you are deploying APM or Edge Gateway, you must have configured the BIG-IP system for DNS and NTP. See *Appendix A: Configuring DNS and NTP settings on the BIG-IP on page 44* for instructions.

### APM

The first section of the iApp in this scenario asks about the BIG-IP Access Policy Manager. You must have APM fully licensed and provisioned to use APM. For more information on BIG-IP APM, see <http://www.f5.com/products/big-ip/access-policy-manager.html>.

1. **Outlook Web App FQDN**

If you are deploying Outlook Web App, type the fully qualified domain name (FQDN) that clients use to access Outlook Web App. This is the FQDN you must configure in DNS to resolve to the IP address of the BIG-IP virtual server providing services for Outlook Web App (OWA). Do not include the protocol (e.g. "https://") or any path information (e.g. "/owa"). If you are not deploying OWA at this time, you may leave this field blank.

**Note**

*While the iApp only asks for the OWA FQDN, the BIG-IP APM supports all Exchange HTTP-based Client Access Services.*

2. **Active Directory name or IP address**

Type the name or IP address of an Active Directory server in your domain that this BIG-IP system can contact. If you provide a name rather than an IP address, use the FQDN of the server rather than the NetBIOS name. Make sure this BIG-IP system and the Active Directory server have routes to one another and that firewalls allow traffic between the two.

3. **Active Directory domain name**

Type the Active Directory domain name in FQDN format (such as "mydomain.example.com").

4. **Active Directory Anonymous binding**

If your Active Directory implementation does not allow anonymous binding, select **Credentials are required for binding**. Two new questions appear.

i. *Active Directory user name*

Type a user name with administrative permissions.

ii. *Password*

Type the associated password.

**Note**

*These credentials are stored in plain text on your BIG-IP system.*

If anonymous binding is allowed, no additional information is required.

This completes the BIG-IP APM section.

## Tell us about your Edge Gateway deployment

This section of the iApp asks about your Edge Gateway deployment.

1. **Edge Gateway virtual server IP address**

Type the IP address you want to use for the Edge Gateway virtual server. This is the address clients will use to access the HTTP-based Client Access services.

2. **SSL Certificate**

Select the SSL certificate you imported onto the BIG-IP system for decrypting client connections. This Edge Gateway/APM deployment supports the use of a single virtual server with one IP address, so if you are using different DNS names for different services, you need to provide a Subject Alternative Name (SAN) enabled certificate and key that is valid for each name.

For detailed information on SAN certificates, see *Subject Alternative Name (SAN) SSL Certificates on page 70*.

3. **SSL Key**

Select the associated key from the list.

4. **Re-encrypt traffic?**

This question asks if you want to re-encrypt the traffic that is forward to the BIG-IP LTM virtual server.

**Important**



*If you are deploying Edge Gateway and forwarding traffic to a separate BIG-IP LTM, we generally recommend that you do not re-encrypt traffic between your BIG-IP Edge Gateway and BIG-IP LTM because both BIG-IP systems must process the SSL transactions. However, if you do choose to re-encrypt, we strongly recommend you use a valid certificate (usually SAN-enabled) rather than the default, self-signed certificate for the Client SSL profile on your BIG-IP LTM. If not re-encrypting traffic, you do not need a certificate on your BIG-IP LTM.*

A. **Re-encrypt (SSL Bridging)**

If you choose **Re-encrypt (SSL Bridging)**, the BIG-IP system encrypts the traffic headed for the BIG-IP LTM virtual server.

- **Remote LTM using self-signed certificate?**

If you choose to Re-encrypt, an additional question appears asking if the remote BIG-IP LTM to which you are sending this traffic is using a BIG-IP self-signed (or default) certificate for decryption, or a certificate signed by a Certificate Authority. This BIG-IP system will not trust the remote BIG-IP default or a self-signed certificate unless specifically configured to do so in this question. Choose the appropriate option from the list.

**Important**



*This question specifically pertains to the certificate used by the remote BIG-IP system, NOT the certificates present and assigned on this local BIG-IP system.*

B. **Do not re-encrypt (SSL Offload)**

If you choose **Do not re-encrypt (SSL Offload)**, the BIG-IP system does not modify the traffic, and you can continue with the next question.

5. ***Virtual server address on the remote BIG-IP LTM***

Type the IP address of the remote BIG-IP LTM virtual server to which you will be forwarding Client Access traffic from this BIG-IP device.

6. ***WAN or LAN***

Choose whether most clients are connecting over a WAN or LAN. This setting is used to configure the proper TCP optimization settings.

7. **URI for reaching Outlook Web App**

If you are deploying Outlook Web App, you must specify the URI for reaching OWA. By default, the URI is **/owa/**. If your URI is different, type it here.

**Note**



*We do not recommend changing the default URI for Outlook Web App. If you modified the URI to use the root directory (**http(s)://<fqdn>/**), you must NOT change this value from the default. .*

### Additional Steps

Review the information in the Additional steps section, and take appropriate action if necessary. All of the notes in Additional Steps are found in the relevant section of this deployment guide.

### Finished

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

**Critical**



*There are important modifications you must make to the configuration produced by the iApp template. See *Modifying the iApp template configuration* on page 35.*

## Modifying the iApp template configuration

This section contains modifications to the configuration produced by the template. These are not required changes in all cases, check to see if the modifications apply to your deployment.

### Disabling Strict Updates

Some modifications require you to disable the Strict Updates feature. If a modification requires disabling Strict Updates, there is a note in the applicable procedure. You only need to turn off the Strict Updates feature once, unless you have re-enabled it.

By disabling Strict Updates, if you re-enter the iApp template and modify the configuration within the iApp, you must make all of the following changes again manually.

#### To disable Strict Updates

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. Click the name of your Exchange 2010 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.
6. You must restart bigd after making changes that require disabling Strict Updates. However, you should only restart bigd after making all necessary modifications in this section. See *Restart bigd after making changes that require disabling Strict Updates on page 42*. This note is repeated in all relevant procedures.

**Critical**



### Required: Modify the Cookie persistence profile timeout value

The iApp template incorrectly sets the Cookie persistence timeout value to 180 seconds. The correct setting should be 0 seconds, which marks the cookie as a session cookie. You must have command line access on the BIG-IP system to modify the timeout value.

#### To modify the Cookie persistence timeout value

1. Disable Strict Updates as described in the procedure on this page.
2. Open a command prompt on the BIG-IP system.
3. If necessary, enter the TMSH shell by typing **tmssh**.
4. At the tmos prompt, use the following command syntax:

```
modify /ltm persistence cookie <app_name>/<profile_name> timeout 0
```

Where <app\_name> is the name you gave the iApp, and profile name is the name of the Cookie persistence profile created by the template. This name is preceded by the name you gave the iApp, followed by **\_cookie\_persistence\_profile**.

In our example, we type **modify /ltm persistence cookie my\_Exchange2010.app/my\_Exchange2010\_cookie\_persistence\_profile timeout 0**

5. You must restart bigd after making changes that require disabling Strict Updates. However, you should only restart bigd after making all necessary modifications in this section. See *Restart bigd after making changes that require disabling Strict Updates on page 42*.

### Required: Check the OneConnect profile

In some scenarios, the template currently leaves the Source Mask of the OneConnect profile at the default (0.0.0.0). This Source Mask should be 255.255.255.255. We recommend checking the Source Mask of your OneConnect profile, and modifying it if necessary.

#### To check the OneConnect profile

1. On the Main tab, expand **Local Traffic** and then click **Profiles**.
2. On the Menu bar, from the **Other** menu, select **OneConnect**.
3. Click the name of the OneConnect profile created by the iApp template. This profile is preceded by the name you gave the iApp, followed by **\_oneconnect**
4. Check the **Source Mask** box.
  - If the Source Mask is **255.255.255.255**, no modifications are necessary.
  - If the Source Mask is **0.0.0.0**, you must first disable Strict Updates as described in *Disabling Strict Updates on page 35*. Once you have disabled Strict Updates, return to this OneConnect profile, check the Custom box for **Source Mask**, and then type **255.255.255.255** in the box. When you are finished, click the **Update** button.
5. You must restart bigd after making changes that require disabling Strict Updates. However, you should only restart bigd after making all necessary modifications in this section. See *Restart bigd after making changes that require disabling Strict Updates on page 42*.

### Required: Modify the iRule(s)

You must modify the iRule(s) produced by the iApp for the following scenarios:

- **Persistence iRule**
  - » You configured the iApp to use a single virtual server for all HTTP-based Client Access services, or
  - » You configured the iApp to use separate virtual servers **and** are deploying Outlook Anywhere
- **Edge Gateway**

You configured the iApp to use Edge Gateway in any scenario. You may need to also modify the Persistence iRule if either of the scenarios described above applies to your implementation.

You must modify the HTTP Response section at the bottom of all iRules that apply to your implementation. There is an additional step (6) for the Persistence iRule using a single virtual server.

#### To modify the iRules

1. If you have not already disabled Strict Updates, follow the instructions in *Disabling Strict Updates on page 35*.
2. On the Main tab, expand **Local Traffic** and then click **iRules**.
3. Click the name of the applicable iRule:
  - **Persistence iRule, single virtual server**

This iRule is preceded by the name you gave the iApp template, followed by **\_combined\_vs\_persist\_iRule**. For example *Exchange2010\_combined\_vs\_persist\_iRule*.

- **Persistence iRule, separate virtual servers and deploying Outlook Anywhere**

This iRule is preceded by the name you have the iApp template, followed by `_oa_persist_iRule`. For example `Exchange2010_oa_persist_iRule`.

- **Edge Gateway**

This iRule is preceded by the name you have the iApp template, followed by `_edge_base_iRule`. For example `Exchange2010_edge_base_iRule`.

4. In the **Definition** section, find the **HTTP Response** statement. If you have not modified the iRule produced by the iApp previously, this statement looks like the following example:

```
when HTTP_RESPONSE {
    if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate" } {
        ONECONNECT::reuse disable
    }
}
```

5. Replace the entire HTTP Response statement with the following statement, omitting the line numbers:

```
1 when HTTP_RESPONSE {
2     if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate" } {
3         ONECONNECT::reuse disable
4         ONECONNECT::detach disable
5         ## this disables NTLM conn pool for connections where OneConnect has been disabled
6         NTLM::disable
7     }
8     ## this command rechunks encoded responses
9     if {[HTTP::header exists "Transfer-Encoding"]} {
10        HTTP::payload rechunk
11    }
12 }
```

6. **Persistence iRule, single virtual server only**

If you configured a single virtual server, you must also add the following lines to the persistence iRule, immediately after the pool statement in the **microsoft-server-activesync**, **ews**, and **/rpc/rpcproxy.dll** sections.

```
COMPRESS::disable
CACHE::disable
```

For example, the ActiveSync section should look like the following:

```
"/microsoft-server-activesync" {
    pool my_Exchange_2010__single_as_pool
    COMPRESS::disable
    CACHE::disable
    persist uie $sessionid 7200
    return
}
```

Be sure to add these two lines to each of the sections.

See the complete iRule in the manual configuration section on [page 54](#) for example syntax.

7. Click the **Update** button.
8. If necessary, repeat this entire procedure for additional iRules as described in step 3.
9. You must restart bigd after making changes that require disabling Strict Updates. However, you should only restart bigd after making all necessary modifications in this section. See *Restart bigd after making changes that require disabling Strict Updates on page 42*.

### Required: New Autodiscover monitor script file

If you configured the iApp for Autodiscover and selected to use advanced monitors, you must download and import a new script file. There are now two different script files, depending on whether you configured SSL Offload or SSL Bridging. A portion of this procedure requires command line access to the BIG-IP system.

#### Important



*This modification is only necessary if you deployed Autodiscover and chose Advanced Monitors*

#### To download and install the script

1. Download the appropriate script:

**If the BIG-IP system is configured for SSL Offload:**

[www.f5.com/solution-center/deployment-guides/files/autodiscover-monitor.zip](http://www.f5.com/solution-center/deployment-guides/files/autodiscover-monitor.zip)

**If the BIG-IP system is configured for SSL Bridging (SSL re-encryption):**

[www.f5.com/solution-center/deployment-guides/files/autodiscover-monitor-ssl-bridging.zip](http://www.f5.com/solution-center/deployment-guides/files/autodiscover-monitor-ssl-bridging.zip)

2. Extract the file to a location accessible by the BIG-IP system.
3. From the Main tab of the BIG-IP Configuration utility, expand **System**, and then click **File Management**.
4. On the Menu bar, click **External Monitor Program File List**.
5. Click the **Import** button.
6. In the **File Name** row, click **Browse**, and then locate the **autodiscover-monitor.sh** (SSL offload) or **autodiscover-monitor-ssl-bridging.sh** (SSL bridging) file.
7. Click the **Overwrite Existing** button, and then from the list, select **/Common/autodiscover\_eav**.
8. Click the **Import** button.

### Required: Modifying the RPC Client Access persistence profile

If you deployed the iApp template for RPC Client Access, you must modify the fallback persistence profile to enable Match Across Virtual Servers. You must also modify either the single virtual server you created for the HTTP-based CAS services, or the separate virtual server you created for Outlook Anywhere to use this persistence profile as a Fallback persistence profile.

#### To modify the RPC Client Access persistence profile

1. If you have not already disabled Strict Updates, follow the instructions in *Disabling Strict Updates on page 35*.
2. On the Main tab, expand **Local Traffic** and then click **Profiles**.

3. On the Menu bar, click **Persistence**.
4. Click the name of the Source Address persistence profile created by the iApp template. This profile is preceded by the name you gave the iApp, followed by **\_source\_address\_persistence\_profile**.
5. Click the **Match Across Virtual Servers** box to enable Match Across Virtual Servers.
6. Click the **Update** button.
7. You must restart bigd after making changes that require disabling Strict Updates. However, you should only restart bigd after making all necessary modifications in this section. See *Restart bigd after making changes that require disabling Strict Updates on page 42*.

The next task is to add this persistence profile as a fallback persistence method to the single virtual server you created for HTTP-based CAS services, or the separate virtual server you created for Outlook Anywhere. Strict Updates must remain disabled.

#### **To modify the virtual servers to use the persistence profile as a fallback**

1. On the Main tab, expand **Local Traffic** and then click **Virtual Servers**.
2. Click the name of the appropriate virtual server:
  - If you created a single virtual server for all HTTP-based CAS services, this virtual server is preceded by the name you gave the iApp, followed by **\_combined\_https**.
  - If you created separate virtual servers and deployed the iApp for Outlook Anywhere, this virtual server is preceded by the name you gave the iApp, followed by **\_oa\_https**.
3. On the Menu bar, click **Resources**.
4. From the **Fallback Persistence Profile** list, select the RPC Source Address persistence profile you just modified.
5. Click **Update**.
6. You must restart bigd after making changes that require disabling Strict Updates. However, you should only restart bigd after making all necessary modifications in this section. See *Restart bigd after making changes that require disabling Strict Updates on page 42*.

#### **Required: Adding the Append iRule to an Edge Gateway configuration**

If you used the iApp template on Edge Gateway (or configured the Edge Gateway manually using the manual configuration tables), and the Edge Gateway is forwarding directly to a BIG-IP LTM that has not been properly configured using either the iApp template or the manual configuration, to a single Exchange 2010 CAS server rather than to a BIG-IP LTM, or to a non-F5 device (the last scenario is unsupported), users may not receive a “friendly” HTTP redirect to ‘FQDN/owa’ when they attempt to access OWA just using the FQDN. You must add the **owa append** iRule to the virtual server on the Edge Gateway system.

#### **Important**



*This is only applicable to the Edge Gateway scenarios mentioned above. Do not follow these steps if you are only using BIG-IP LTM, have combined BIG-IP LTM and APM on a single system, or are using Edge Gateway and forwarding to a BIG-IP LTM system that is configured with the ‘owa append’ iRule.*

### To add the OWA Append 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 unique name such as **edge-owa-append-irule**.
4. In the **Definition** section, copy and paste the following iRule:

```
1  when HTTP_REQUEST {
2      if {[HTTP::uri] == "/" } {
3          HTTP::uri /owa
4      }
5  }
```

5. Click **Finished**.
6. If you configured the Edge Gateway using the iApp, you must disable Strict Updates. Use the procedure *Disabling Strict Updates on page 35*.
7. On the Main tab, click **Virtual Servers**.
8. From the Virtual Server list, click the name of the Edge Gateway virtual server. If you configured the Edge Gateway using the iApp, this name is prefaced by the name you gave the iApp, followed by **\_edge**.
9. On the Menu bar, click **Resources**.
10. From the iRules section, click **Manage**.
11. From the **Available** list, select the append iRule you just created and then click the Add (<<) button.
12. Click the **Up** button to move the Append iRule to the top of the list.
13. Click **Finished**.
14. You must restart bigd after making changes that require disabling Strict Updates. However, you should only restart bigd after making all modifications in this section. See *Restart bigd after making changes that require disabling Strict Updates on page 42*.

### Optional: Adding a user name and password to POP3 and IMAP4 simple monitors

If you chose to use simple monitors **and** to configure the iApp to handle POP3 or IMAP4, or both, you must add a user name and password to the monitor(s) after completing the iApp. You can add the User Name and Password values without disabling Strict Updates.

#### Important



*This modification is only necessary if you are using simple health monitors and POP3 and/or IMAP4.*

### To modify the POP3 and IMAP4 simple monitors to include a user name and password

1. On the Main tab, expand **Local Traffic** and then click **Monitors**.
2. From the Monitor list, click the name of the POP3 or IMAP4 monitor created by the iApp. The monitor is prefaced by the name you gave the iApp, followed by **\_simple\_pop3\_monitor** or **\_simple\_imap4\_monitor**.

3. In the **User Name** box, type a valid user name with a POP3 or IMAP4 account. We strongly recommend creating user accounts specifically for these monitors.
4. In the **Password** box, type the associated password.
5. Click the **Update** button.

### Optional: Additional health monitors for RPC Client Access

F5 strongly recommends using Outlook Anywhere (RPC over HTTP) rather than RPC Client Access (MAPI). Outlook Anywhere can be monitored more accurately, and is more likely to work through firewalls because it uses HTTPS exclusively for connectivity. Additionally, you can elect to offload SSL processing from your Client Access Servers; cryptographic operations for RPC Client Access, which uses a different encryption method, cannot be offloaded.

If you chose to implement RPC Client Access, you may want to implement additional monitoring steps beyond the basic monitor provided by the iApp template.

To create the health monitors, from the Main tab of the BIG-IP Configuration utility, expand **Local Traffic**, and then click **Monitors**. Click the **Create** button, and then use the following table to configure the health monitor options. The table contains any non-default settings required or recommended for this configuration.

Monitor Field	Description/Notes
<b>MAPI Monitor</b>	
<b>Name</b>	Type a meaningful name, such as <b>MAPI_59532_TCP_monitor</b>
<b>Type</b>	<b>TCP</b>
<b>Interval</b>	User choice, but we recommend <b>30</b>
<b>Timeout</b>	User choice, but we recommend <b>91</b>
<b>Alias Service Port<sup>1</sup></b>	Type the port that matches the RPC Client Access static port for MAPI on your Client Access Servers. In our example, we use <b>59532</b> .
<b>Address Book Monitor</b>	
<b>Name</b>	Type a meaningful name, such as <b>AddressBook_59533_TCP_monitor</b>
<b>Type</b>	<b>TCP</b>
<b>Interval</b>	User choice, but we recommend <b>30</b>
<b>Timeout</b>	User choice, but we recommend <b>91</b>
<b>Alias Service Port<sup>1</sup></b>	Type the port that matches the RPC Client Access static port for the Address Book on your Client Access Servers. In our example, we use <b>59533</b> .

<sup>1</sup> You must select **Advanced** from the **Configuration** list for this option to appear

### Attaching the monitors to the RPC Client Access pool

After you have created the health monitors, the final task in this section is to attach the new monitors to the RPC Client Access pool.

#### To add the monitors to the pool

1. If you have not already disabled Strict Updates, follow the instructions in *Disabling Strict Updates* on page 35.
2. On the Main tab, expand **Local Traffic** and then click **Pools**.
3. Click the name of the RPC Client Access pool created by the template. This pool name is

preceded by the name you have the iApp template, followed by **\_rpc\_pool**.

4. In the Health Monitor section, from the **Available** list, select one of the monitors you just created and then click the Add (<<) button. Repeat for the other monitor. Note the existing monitor should remain in the **Active** box.
5. Click the **Update** button.
6. You must restart bigd after making changes that require disabling Strict Updates. However, you should only restart bigd after making all modifications in this section. See the following section.

### Restart bigd after making changes that require disabling Strict Updates

After performing any modification that requires disabling Strict Updates feature on the Application Service, you must restart the bigd daemon from the BIG-IP command line. We recommend restarting bigd during a maintenance window or other scheduled downtime.

#### To restart bigd

1. From the command line, log into the BIG-IP system.
2. From the prompt, run the following command:

```
bigstart restart bigd
```

3. Exit the command line interface.

This completes the modifications.

## Next steps

After completing the iApp Template, the BIG-IP Application Services page opens for the Exchange 2010 service you just created. To see the list of all the configuration objects created to support Exchange 2010, on the Menu bar, click **Components**. The complete list of all Exchange 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 Exchange 2010 implementation to point to the BIG-IP system's virtual server address.

## Viewing statistics

You can easily view a number of different statistics on the BIG-IP system related to the Exchange 2010 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 Exchange 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 Exchange 2010 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 Exchange 2010 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 A: Configuring DNS and NTP settings on the BIG-IP

If you are configuring the iApp to use BIG-IP Edge Gateway or APM, you must configure DNS and NTP settings on the BIG-IP system before beginning the iApp.

### Configuring the DNS settings

In this section, you configure the DNS settings on the BIG-IP to point to the Active Directory server.

- **Note:** *DNS lookups go out over one of the interfaces configured on the BIG-IP system, not the management interface. The management interface has its own, separate DNS settings.*
  
- **Important:** *The BIG-IP system must have a Route to the Active Directory server. The Route configuration is found on the Main tab by expanding **Network** and then clicking **Routes**. For specific instructions on configuring a Route on the BIG-IP system, see the online help or the product documentation.*

#### To configure DNS settings

1. On the Main tab, expand **System**, and then click **Configuration**.
2. On the Menu bar, from the **Device** menu, click **DNS**.
3. In the **DNS Lookup Server List** row, complete the following:
  - A. In the **Address** box, type the IP address of a DNS server that can resolve the Active Directory server.
  - B. Click the **Add** button.
4. Click **Update**.

### Configuring the NTP settings

The next task is to configure the NTP settings on the BIG-IP system for authentication to work properly.

#### To configure NTP settings

1. On the Main tab, expand **System**, and then click **Configuration**.
2. On the Menu bar, from the **Device** menu, click **NTP**.
3. In the **Address** box, type the fully-qualified domain name (or the IP address) of the time server that you want to add to the Address List.
4. Click the **Add** button.
5. Click **Update**.

To verify the NTP setting configuration, you can use the **ntpq** utility. From the command line, run **ntpq -np**.

See <http://support.f5.com/kb/en-us/solutions/public/10000/200/sol10240.html> for more information on this command.

## Appendix B: Using X-Forwarded-For to log the client IP address in IIS 7.0 and 7.5 (optional)

When you configure BIG-IP LTM to use SNAT, the BIG-IP system replaces the source IP address of an incoming connection with its local self IP address (in the case of SNAT Automap), or an address you have configured in a SNAT pool. As a result, Microsoft IIS logs each connection with its assigned SNAT address, rather than the address of the client. The iApp produces an HTTP profile on the BIG-IP system which inserts an X-Forwarded-For header, so the original client IP address is sent as well; however, in default IIS configuration, this information is not logged.

Beginning with IIS 7, Microsoft provides an optional Advanced Logging Feature for IIS that allows you to define custom log definitions that can capture additional information such as the client IP address included in the X-Forwarded-For header.

### Deploying the Custom Logging role service

The first task is to deploy the Custom Logging role service. If you do not deploy this role service, you may receive a "Feature not supported" error when trying to edit the log definition in the next section.

#### To deploy the Custom Logging role service

1. From your Windows Server 2008 or Windows Server 2008 R2 device, open Server Manager.
2. In the Navigation pane, expand **Roles**.
3. Right-click **Web Server**, and then click **Add Role Services**.
4. Under Health and Diagnostics, check the box for **Custom Logging**, and then click **Next**.
5. On the Confirmation page, click **Install**.
6. After the service has successfully installed, click the **Close** button.

### Adding the X-Forwarded-For log field to IIS

Before beginning the following procedure, you must have installed IIS Advanced Logging. For installation instructions, see

[http://www.iis.net/community/files/media/advancedlogging\\_readme.htm](http://www.iis.net/community/files/media/advancedlogging_readme.htm)

If you are using IIS version 6, F5 has a downloadable ISAPI filter that performs a similar function to the Advanced Logging Feature discussed here. For information on that solution, see the DevCentral post at [http://devcentral.f5.com/weblogs/Joel/archive/2009/08/19/x\\_forwarded\\_for\\_log\\_filter\\_for\\_windows\\_servers.aspx](http://devcentral.f5.com/weblogs/Joel/archive/2009/08/19/x_forwarded_for_log_filter_for_windows_servers.aspx)

#### To add the X-Forwarded-For log field to IIS

1. From your Windows Server 2008 or Windows Server 2008 R2 device, open the Internet Information Services (IIS) Manager.
2. From the Connections navigation pane, click the appropriate server, web site, or directory on which you are configuring Advanced Logging. The Home page appears in the main panel.
3. From the Home page, under IIS, double-click **Advanced Logging**.
4. From the Actions pane on the right, click **Edit Logging Fields**.

5. From the Edit Logging Fields dialog box, click the **Add Field** button, and then complete the following:
  - A. In the **Field ID** box, type **X-Forwarded-For**.
  - B. From the **Category** list, select **Default**.
  - C. From the **Source Type** list, select **Request Header**.
  - D. In the **Source Name** box, type **X-Forwarded-For**.
  - E. Click the **OK** button.
6. Click a Log Definition to select it. By default, there is only one: %COMPUTERNAME%-Server. The log definition you select must have a status of Enabled.
7. From the Actions pane on the right, click **Edit Log Definition**.
8. Click the **Select Fields** button, and then check the box for the X-Forwarded-For logging field.
9. Click the **OK** button.
10. From the Actions pane, click **Apply**.
11. Click **Return To Advanced Logging**.
12. In the Actions pane, click **Enable Advanced Logging**.

Now, when you look at the logs, the client IP address is included.

## Appendix C: Manual configuration tables

This table contains the BIG-IP configuration objects in this deployment and any non-default settings for advanced users. See the Edge Gateway and APM tables for additional Edge Gateway and APM configuration. Give each BIG-IP object a unique name in the **Name** field.

### Configuration table if using a single virtual server for Exchange HTTP-based services

BIG-IP object	Non-default settings/Notes			
<b>Health Monitors</b>	<b>Outlook Web App</b>	HTTP parent (see <i>Outlook Web App advanced monitor on page 58 for optional, recommended monitor</i> )		
	<b>Outlook Anywhere</b>	HTTP parent (see <i>Outlook Anywhere advanced monitor on page 58 for optional, recommended monitor</i> )		
	<b>ActiveSync</b>	HTTP parent (see <i>ActiveSync advanced monitor on page 59 for optional, recommended monitor</i> )		
	<b>Autodiscover</b>	HTTP parent (see <i>Autodiscover advanced EAV monitor on page 59 for optional, recommended monitor</i> )		
	<b>Important note:</b> If configuring SSL Bridging only, use the <b>HTTPS</b> parent for all monitors			
<b>Pools</b> (repeat for each Client Access Server role)	<b>Health monitor</b>	Add the appropriate health monitor for the Client Access role you created above		
	<b>Slow Ramp Time</b>	<b>300</b> (must select Advanced from the Configuration menu for this option to appear)		
	<b>Load Balancing Method</b>	<b>Least Connections (member)</b> recommended		
	<b>Address</b>	IP Address of Client Access server running Outlook Web App		
	<b>Service Port</b>	<b>80 (443</b> if configuring SSL Bridging) Repeat Address and Port for all members <b>Important:</b> Create a pool for each Client Access Server role		
<b>iRules</b>	Append (page 54), Persistence iRule (page 54). <b>Important:</b> The Append iRule should be listed first when configuring the virtual server			
<b>Profiles</b>	<b>HTTP</b>	Parent Profile	<b>http</b>	
		Redirect Rewrite	<b>All</b>	
	<b>HTTP Compression</b>	Content List-->Include List (Add each entry to the Content Type box and click Include. This is optional but recommended.)		application/vnd.ms-publisher
				application/(xls excel msexcel ms-excel x-excel x-xls xmsexcel x-ms-excel vnd.excel vnd.msexcel vnd.ms-excel)
				application/(word doc msword winword ms-word x-word x-msword vnd.word vnd.msword vnd.ms-word)
				application/(xml x-javascript javascript x-ecmascript ecmascript)
				application/(powerpoint msppowerpoint ms-powerpoint x-powerpoint x-mspowerpoint vnd.powerpoint vnd.mspowerpoint  vnd.ms-powerpoint vnd.ms-pps)
				application/(mpp msproject x-msproject x-ms-project vnd.ms-project)
				application/(visio x-visio vnd.visio vsd x-vsd x-vsd)
				application/(pdf x-pdf acrobat vnd.pdf)
<b>Web Acceleration</b>	Parent Profile	<b>optimized-caching</b>		
<b>TCP WAN<sup>1</sup></b>	Parent Profile	<b>tcp-wan-optimized</b>		
<b>TCP LAN<sup>1</sup></b>	Parent Profile	<b>tcp-lan-optimized</b>		
<b>Client SSL</b>	Parent Profile	<b>clientssl</b>		
	Certificate/Key	Select the Certificate and Key you imported		
<b>Server SSL<sup>2</sup></b>	Parent Profile	<b>serverssl</b>		
<b>Persistence<sup>3</sup></b>	Persistence Type	<b>Cookie<sup>3</sup></b>		
<b>OneConnect</b>	Parent Profile	<b>oneconnect</b>		
	Source Mask	<b>255.255.255.255</b>		
<b>NTLM</b>	Parent Profile	<b>ntlm</b>		
<b>Virtual Servers</b>	<b>Port 443</b>	Destination Address	IP address for the virtual server (Service Port <b>443</b> )	
		Profiles	Add each of the profiles you created above from the appropriate list	
		SNAT Pool	<b>Automap<sup>4</sup></b>	
		iRules	Append, Persistence (the Append iRule must be listed first)	
	Default Pool	Do <b>not</b> select a default pool for this virtual		
<b>Port 80</b> (optional, for redirect purposes only)	Destination Address	IP address for the virtual server (Service Port <b>80</b> )		
	Profiles	HTTP profile only		
	iRule	<b>_sys_https_redirect</b>		

<sup>1</sup> The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent

<sup>2</sup> Server SSL profile is only necessary if configuring SSL Bridging.

<sup>3</sup> **Important:** See Required: Modify the Cookie persistence profile timeout value on page 35 for an important modification to this profile.

<sup>4</sup> If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Automap. See the BIG-IP documentation for creating SNAT Pools.

## Configuration table if using separate virtual servers for Exchange HTTP-based services

### Outlook Web App configuration table - includes the Exchange Control Panel (ECP)

BIG-IP object	Non-default settings/Notes		
<b>Health Monitor</b>	<b>Type</b>	<b>HTTP</b> parent (use <b>HTTPS</b> parent if configuring SSL Bridging) For recommended, optional monitor, see <i>Outlook Web App advanced monitor on page 58</i>	
<b>Pool</b>	<b>Health monitor</b>	Add the health monitor you created above	
	<b>Slow Ramp Time</b>	<b>300</b> (must select Advanced from the Configuration menu for this option to appear)	
	<b>Load Balancing Method</b>	<b>Least Connections (member)</b> recommended	
	<b>Address</b>	IP Address of Client Access server running Outlook Web App	
	<b>Service Port</b>	<b>80 (443</b> if configuring SSL Bridging) Repeat Address and Port for all members	
<b>iRules</b>	Append (page 54), Persistence iRule (page 54). <b>Important:</b> The Append iRule should be listed first when configuring the virtual server		
<b>Profiles</b>	<b>HTTP</b>	Parent Profile Redirect Rewrite <b>http</b> <b>All</b>	
	<b>HTTP Compression</b>	Content List-->Include List (Add each entry to the Content Type box and click Include. This is optional but recommended.)	application/vnd.ms-publisher
			application/(xls excel msexcel ms-excel x-excel x-xls x-msexcel x-ms-excel vnd.excel vnd.msexcel vnd.ms-excel)
			application/(word doc msword winword ms-word x-word x-msword vnd.word vnd.msword vnd.ms-word)
			application/(xml x-javascript javascript x-ecmascript ecmascript)
			application/(powerpoint mspowerpoint ms-powerpoint x-powerpoint x-mspowerpoint vnd.powerpoint vnd.mspowerpoint  vnd.ms-powerpoint vnd.ms-pps)
			application/(mpp msproject x-msproject x-ms-project vnd.ms-project)
			application/(visio x-visio vnd.visio vsd x-vsd x-vsd)
	application/(pdf x-pdf acrobat vnd.pdf)		
	<b>Web Acceleration</b>	Parent Profile <b>optimized-caching</b>	
<b>TCP WAN<sup>1</sup></b>	Parent Profile <b>tcp-wan-optimized</b>		
<b>TCP LAN<sup>1</sup></b>	Parent Profile <b>tcp-lan-optimized</b>		
<b>Client SSL</b>	Parent Profile Certificate/Key <b>clientssl</b> Select the Certificate and Key you imported		
<b>Server SSL<sup>2</sup></b>	Parent Profile <b>serverssl</b>		
<b>Persistence<sup>3</sup></b>	Persistence Type <b>Cookie<sup>3</sup></b>		
<b>OneConnect</b>	Parent Profile Source Mask <b>oneconnect</b> <b>255.255.255.255</b>		
<b>NTLM</b>	Parent Profile <b>ntlm</b>		
<b>Virtual Servers</b>	<b>Port 443</b>	Destination Address	IP address for the virtual server (Service Port <b>443</b> )
		Profiles	Add each of the profiles you created above from the appropriate list
SNAT Pool		Automap <sup>4</sup>	
iRules		Append, Persistence (the Append iRule must be listed first)	
Default Pool		Select the pool you created for Outlook Web App above	
<b>Port 80</b> (optional, for redirect purposes only)	Destination Address	IP address for the virtual server (Service Port <b>80</b> )	
	Profiles	HTTP profile only	
	iRule	_sys_https_redirect	

<sup>1</sup> The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent

<sup>2</sup> Server SSL profile is only necessary if configuring SSL Bridging

<sup>3</sup> **Important:** See Required: Modify the Cookie persistence profile timeout value on page 35 for an important modification to this profile.

<sup>4</sup> If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Automap. See the BIG-IP documentation for creating SNAT Pools.

[Outlook Anywhere configuration table \(for separate virtual servers\) - includes EWS \(Exchange Web Services\) and OAB \(Offline Address Book\)](#)

BIG-IP object	Non-default settings/Notes		
<b>Health Monitor</b>	<b>Type</b>	<b>HTTP</b> (use <b>HTTPS</b> parent if configuring SSL Bridging) For recommended, optional monitor, see <i>Outlook Anywhere advanced monitor on page 58</i>	
<b>Pool</b>	<b>Health monitor</b>	Add health monitor above	
	<b>Slow Ramp Time<sup>1</sup></b>	<b>300</b>	
	<b>Load Balancing Method</b>	<b>Least Connections (member)</b> recommended	
	<b>Address</b>	IP Address of Client Access server running Outlook Anywhere	
	<b>Service Port</b>	<b>80 (443</b> if configuring SSL Bridging) Repeat Address and Port for all members	
<b>iRules</b>	<b>OA Persist</b> (page 57) which is associated with the persistence profile below and the built-in iRule: <b>_sys_http_redirect<sup>1</sup></b> if offloading SSL		
<b>Profiles</b>	<b>HTTP</b>	Parent Profile	<b>http</b>
		<b>Redirect Rewrite</b>	<b>Match</b>
	<b>TCP WAN<sup>2</sup></b>	Parent Profile	<b>tcp-wan-optimized</b>
	<b>TCP LAN<sup>2</sup></b>	Parent Profile	<b>tcp-lan-optimized</b>
	<b>Client SSL</b>	Parent Profile	<b>clientssl</b>
		Certificate/Key	Select the Certificate and Key you imported
	<b>Server SSL<sup>3</sup></b>	Parent Profile	<b>serverssl</b>
	<b>OneConnect</b>	Parent Profile	<b>oneconnect</b>
Source Mask		<b>255.255.255.255</b>	
<b>NTLM</b>	Parent Profile	<b>ntlm</b>	
<b>Persistence</b>	Persistence Type	<b>Universal</b>	
	iRule	Select the OA Persist iRule you created above	
<b>Virtual Servers</b>	<b>Port 443</b>	Destination Address	IP address for the virtual server (Service Port <b>443</b> )
		Profiles	Add each of the profiles you created above from the appropriate list
		SNAT Pool	<b>Automap<sup>4</sup></b>
	Default Pool	Select the pool you created for Outlook Anywhere above	
<b>Port 80</b> (optional, for redirect purposes only)	Destination Address	IP address for the virtual server (Service Port <b>80</b> )	
	Profiles	HTTP profile only	
	iRule	<b>_sys_https_redirect</b>	

[Active Sync manual configuration table \(for separate virtual server configuration\)](#)

BIG-IP object	Non-default settings/Notes		
<b>Health Monitor</b>	<b>Type</b>	<b>HTTP</b> (use <b>HTTPS</b> parent if configuring SSL Bridging) For recommended, optional monitor, see <i>ActiveSync advanced monitor on page 59</i>	
<b>Pool</b>	<b>Health monitor</b>	Add health monitor above	
	<b>Slow Ramp Time<sup>1</sup></b>	<b>300</b>	
	<b>Load Balancing Method</b>	<b>Least Connections (member)</b> recommended	
	<b>Address</b>	IP Address of Client Access server running ActiveSync	
	<b>Service Port</b>	<b>80 (443</b> if configuring SSL Bridging) Repeat Address and Port for all members	
<b>Profiles</b>	<b>HTTP</b>	Parent Profile	<b>http</b>
		<b>Redirect Rewrite</b>	<b>Match</b>
	<b>TCP WAN<sup>2</sup></b>	Parent Profile	<b>tcp-wan-optimized</b>
	<b>TCP LAN<sup>2</sup></b>	Parent Profile	<b>tcp-lan-optimized</b>
	<b>Client SSL</b>	Parent Profile	<b>clientssl</b>
Certificate/Key		Select the Certificate and Key you imported	
<b>Server SSL<sup>3</sup></b>	Parent Profile	<b>serverssl</b>	

<sup>1</sup> You must select Advanced from the Configuration list for this option to appear

<sup>2</sup> The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent

<sup>3</sup> Server SSL profile is only necessary if configuring SSL Bridging.

<sup>4</sup> If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Automap. See the BIG-IP documentation for creating SNAT Pools.

[Active Sync manual configuration table for separate virtual server configuration \(continued\)](#)

BIG-IP object	Non-default settings/Notes		
<b>Virtual Servers</b>	<b>Port 443</b>	Destination Address	IP address for the virtual server (Service Port <b>443</b> )
		Profiles	Add each of the profiles you created above from the appropriate list
		SNAT Pool	<b>Automap<sup>4</sup></b>
		Default Pool	Select the pool you created for ActiveSync above
<b>Port 80</b> (optional, for redirect purposes only)	Destination Address	IP address for the virtual server (Service Port <b>80</b> )	
	Profiles	HTTP profile only	
	iRule	<b>_sys_https_redirect</b>	

[Autodiscover manual configuration table \(for separate virtual server configuration\)](#)

BIG-IP object	Non-default settings/Notes		
<b>Health Monitor</b>	<b>Type</b>	<b>HTTP</b> (use <b>HTTPS</b> parent if configuring SSL Bridging) For recommended, optional monitor, see <i>Autodiscover advanced EAV monitor on page 59</i> )	
<b>Pool</b>	<b>Health monitor</b>	Add health monitor above	
	<b>Slow Ramp Time<sup>1</sup></b>	<b>300</b>	
	<b>Load Balancing Method</b>	<b>Least Connections (member)</b> recommended	
	<b>Address</b>	IP Address of Client Access server running Autodiscover	
	<b>Service Port</b>	<b>80 (443 if configuring SSL Bridging)</b> Repeat Address and Port for all members	
<b>Profiles</b>	<b>HTTP</b>	Parent Profile	<b>http</b>
	<b>TCP WAN<sup>2</sup></b>	Parent Profile	<b>tcp-wan-optimized</b>
	<b>TCP LAN<sup>2</sup></b>	Parent Profile	<b>tcp-lan-optimized</b>
	<b>Client SSL</b>	Parent Profile	<b>clientssl</b>
		Certificate/Key	Select the Certificate and Key you imported
<b>Virtual Servers</b>	<b>Port 443</b>	Destination Address	IP address for the virtual server (Service Port <b>443</b> )
		Profiles	Add each of the profiles you created above from the appropriate list
		SNAT Pool	<b>Automap<sup>4</sup></b>
		Default Pool	Select the pool you created for Autodiscover above
	<b>Port 80</b> (optional, for redirect purposes only)	Destination Address	IP address for the virtual server (Service Port <b>80</b> )
	Profiles	HTTP profile only	
	iRule	<b>_sys_https_redirect</b>	

<sup>1</sup> You must select Advanced from the Configuration list for this option to appear

<sup>2</sup> The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent

<sup>3</sup> Server SSL profile is only necessary if configuring SSL Bridging.

<sup>4</sup> If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Automap. See the BIG-IP documentation for creating SNAT Pools.

## Configuration tables for RPC Client Access, POP3, and IMAP4

Use the following tables for RPC Client Access, POP3, and IMAP4, no matter which HTTP-based configuration you chose in the tables on the previous pages. For RPC Client Access, you must decide whether you will use static ports or the default dynamic port range for RPC Client Access traffic. Use the table appropriate for your configuration.

Note that if deploying RPC Client Access, you must also deploy Outlook Anywhere, to properly handle EWS (Exchange Web Services) traffic.

### RPC Client Access<sup>1</sup> dynamic port range manual configuration table

BIG-IP Object	Non-default settings/Notes	
<b>Health Monitor</b>	<b>Type</b>	<b>TCP</b>
	<b>Interval</b>	<b>30</b> (recommended)
	<b>Timeout</b>	<b>91</b> (recommended)
	<b>Alias Service Port</b>	<b>135</b>
<b>Pool</b>	<b>Health monitor</b>	Add health monitor above.
	<b>Action on Service Down<sup>2</sup></b>	<b>Reject</b>
	<b>Slow Ramp Time<sup>2</sup></b>	<b>300</b>
	<b>Load Balancing Method</b>	<b>Least Connections (member)</b> recommended
	<b>Address</b>	IP Address of Client Access server running RPC Client Access
	<b>Service Port</b>	<b>* All Services</b> (repeat Address and Port for all members)
<b>Profiles</b>	<b>Persistence</b>	Parent Profile Timeout Match Across Services Match Across Virtual Servers
		<b>Source Address Affinity</b> <b>7200</b> Click a check in the <b>Match Across Services</b> box Click a check in the <b>Match Across Virtual Servers</b> box
	<b>TCP WAN<sup>3</sup></b>	Parent Profile Idle Timeout
		<b>tcp-wan-optimized</b> <b>7200</b>
	<b>TCP LAN<sup>3</sup></b>	Parent Profile Idle Timeout
		<b>tcp-lan-optimized</b> <b>7200</b>
<b>Virtual Servers</b>	<b>Port 135</b>	<b>Destination Address</b> <b>Service Port</b> <b>Profiles</b> <b>SNAT Pool</b> <b>Default Pool</b>
		IP address for the virtual server <b>135</b> Add each of the profiles you created above from the appropriate list <b>Automap<sup>4</sup></b> Select the pool you created for RPC Client Access above
	<b>All Ports</b>	<b>Destination Address</b> <b>Service Port</b> <b>Profiles</b> <b>SNAT Pool</b> <b>Default Pool</b>
		Same IP address used above (make sure you use a unique name) <b>*All Ports</b> Add each of the profiles you created above from the appropriate list <b>Automap<sup>4</sup></b> Select the pool you created for RPC Client Access above
<b>Additional steps</b>	After completing this virtual server, you must modify either the Single virtual server you created for the HTTP-based CAS services, or the separate virtual server you created for Outlook Anywhere to use the persistence profile you created in this section as a <b>Fallback</b> persistence profile. From the <b>Fallback Persistence Profile</b> list of the Single virtual, or the Outlook Anywhere separate virtual, select the profile you created in this section, and then click the <b>Update</b> button.	

<sup>1</sup> In Exchange Server 2010, you must configure a Client Access Array for your site to use the FQDN you have set to resolve to the IP address of the BIG-IP LTM virtual server, and you must update the existing mailbox database attributes to use that array.

<sup>2</sup> You must select **Advanced** from the Configuration list for this option to appear

<sup>3</sup> The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent, but you must have an Idle Timeout of 7200. The TCP timeout on the BIG-IP is designed to reset idle connections that have become orphaned without a proper close, and gets reset with every packet in a TCP connection. This commonly happens when a client loses network connectivity mid-session. It's good to clear these connections so they don't build up in the connection table.

<sup>4</sup> If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Automap. See the BIG-IP documentation for creating SNAT Pools.

RPC Client Access<sup>1</sup> static ports configuration table

BIG-IP Object	Non-default settings/Notes	
Health Monitors	<b>RPC Monitor</b>	
	<b>Type</b>	TCP
	<b>Interval</b>	30 (recommended)
	<b>Timeout</b>	91 (recommended)
	<b>MAPI Monitor</b>	
	<b>Type</b>	TCP
	<b>Interval</b>	30 (recommended)
	<b>Timeout</b>	91 (recommended)
	<b>Alias Service Port<sup>2</sup></b>	59532 This is the default. Modify this port to match the RPC Client Access static port for MAPI on your Client Access Servers.
	<b>Address Book Monitor</b>	
	<b>Type</b>	TCP
	<b>Interval</b>	30 (recommended)
<b>Timeout</b>	91 (recommended)	
<b>Alias Service Port<sup>2</sup></b>	59533 This is the default. Modify this port to match the RPC Client Access static port for the Address Book on your Client Access Servers.	
Pools	<b>Health monitor</b>	Add health monitor above.
	<b>Action on Service Down<sup>2</sup></b>	Reject
	<b>Slow Ramp Time<sup>2</sup></b>	300
	<b>Load Balancing Method</b>	Least Connections (member) recommended
	<b>Address</b>	IP Address of Client Access server running RPC Client Access
	<b>Service Port</b>	135 (repeat Address and Port for all members)
Create two additional pools, one for <b>MAPI</b> and one for <b>Address Book Service</b> , using the settings above; only the <b>Name</b> , <b>Health Monitor</b> and <b>Service Port</b> are different. Apply the associated Health Monitor you created. The Service Port depends on your configuration.		
Profiles	<b>Persistence</b>	Parent Profile Timeout Match Across Services Match Across Virtual Servers
		<b>Source Address Affinity</b> <b>7200</b> Click a check in the <b>Match Across Services</b> box Click a check in the <b>Match Across Virtual Servers</b> box
	<b>TCP WAN<sup>3</sup></b>	Parent Profile <b>tcp-wan-optimized</b>
	<b>TCP LAN<sup>3</sup></b>	Parent Profile <b>tcp-lan-optimized</b>
Virtual Servers	<b>Destination Address</b>	IP address for the virtual server
	<b>Service Port</b>	135
	<b>Profiles</b>	Add each of the profiles you created above from the appropriate list
	<b>SNAT Pool</b>	Automap <sup>4</sup>
	<b>Default Pool</b>	Select the pool with members using Service Port 135 you created for RPC Client Access above
Create two additional virtual servers, one for <b>MAPI</b> and one for <b>Address Book Service</b> , using the settings above; only the <b>Name</b> , <b>Service Port</b> and <b>Pool</b> are different: The Service Port depends on your configuration. Use the associated pool you created.		
Additional steps	After completing this virtual server, you must modify either the Single virtual server you created for the HTTP-based CAS services, or the separate virtual server you created for Outlook Anywhere to use the persistence profile you created in this section as a <b>Fallback</b> persistence profile. From the <b>Fallback Persistence Profile</b> list of the Single virtual, or the Outlook Anywhere separate virtual, select the profile you created in this section, and then click the <b>Update</b> button.	

<sup>1</sup> In Exchange Server 2010, you must configure a Client Access Array for your site to use the FQDN you have set to resolve to the IP address of the BIG-IP LTM virtual server, and you must update the existing mailbox database attributes to use that array.

<sup>2</sup> You must select Advanced from the Configuration list for this option to appear

<sup>3</sup> The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent

<sup>4</sup> If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Automap. See the BIG-IP documentation for creating SNAT Pools.

### POP3 manual configuration table

BIG-IP Object	Non-default settings/Notes	
<b>Health Monitor</b>	<b>Type</b>	<b>POP3</b> (you must add a User Name and Password of a POP3 user account)
<b>Pool</b>	<b>Health monitor</b>	Add health monitor above
	<b>Slow Ramp Time<sup>2</sup></b>	<b>300</b>
	<b>Load Balancing Method</b>	<b>Least Connections (member)</b> recommended
	<b>Address</b>	IP Address of Client Access server running POP3
	<b>Service Port</b>	<b>110</b> (repeat Address and Port for all members)
<b>Profiles</b>	<b>Client SSL</b>	Parent Profile <b>clientssl</b> Certificate/Key Select the Certificate and Key you imported
	<b>Server SSL<sup>5</sup></b>	Parent Profile <b>serverssl</b>
	<b>TCP WAN<sup>2</sup></b>	Parent Profile <b>tcp-wan-optimized</b>
	<b>TCP LAN<sup>2</sup></b>	Parent Profile <b>tcp-lan-optimized</b>
<b>Virtual Server</b>	<b>Destination Address</b>	IP address for the virtual server
	<b>Service Port</b>	<b>995</b>
	<b>Profiles</b>	Add each of the profiles you created above from the appropriate list
	<b>SNAT Pool</b>	<b>Automap<sup>4</sup></b>
	<b>Default Pool</b>	Select the pool you created for POP3 above

### IMAP4 manual configuration table

BIG-IP Object	Non-default settings/Notes	
<b>Health Monitor</b>	<b>Type</b>	<b>IMAP4</b> (you must add a User Name and Password of a IMAP4 user account)
<b>Pool</b>	<b>Health monitor</b>	Add health monitor above
	<b>Slow Ramp Time<sup>1</sup></b>	<b>300</b>
	<b>Load Balancing Method</b>	<b>Least Connections (member)</b> recommended
	<b>Address</b>	IP Address of Client Access server running IMAP4
	<b>Service Port</b>	<b>143</b> (repeat Address and Port for all members)
<b>Profiles</b>	<b>Client SSL</b>	Parent Profile <b>clientssl</b> Certificate/Key Select the Certificate and Key you imported
	<b>Server SSL<sup>3</sup></b>	Parent Profile <b>serverssl</b>
	<b>TCP WAN<sup>2</sup></b>	Parent Profile <b>tcp-wan-optimized</b>
	<b>TCP LAN<sup>2</sup></b>	Parent Profile <b>tcp-lan-optimized</b>
<b>Virtual Server</b>	Destination Address	IP address for the virtual server
	Service Port	<b>993</b>
	Profiles	Add select each of the profiles you created above from the appropriate list
	SNAT Pool	<b>Automap<sup>4</sup></b>
	Default Pool	Select the pool you created for IMAP4 above

<sup>1</sup> You must select Advanced from the Configuration list for this option to appear

<sup>2</sup> The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent

<sup>3</sup> Server SSL profile is only necessary if configuring SSL Bridging.

<sup>4</sup> If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Automap. See the BIG-IP documentation for creating SNAT Pools.

## iRules

This section contains the iRule code referred to from the manual configuration table. The line numbers are provided for reference. Create a new iRule and copy the code, omitting the line numbers. You may need to modify pool names according to your configuration.

### Append iRule

```
1  when HTTP_REQUEST {  
2      if {[HTTP::uri] == "/" } {  
3          HTTP::uri /owa  
4      }  
5  }
```

#### Important

*This iRule should appear at the top of the iRule list in the virtual server and come before any persistence iRules you might use.*

### Persistence iRule if using a single virtual server for all HTTP-based services

For this configuration, you must create an additional iRule which changes persistence methods based on the service being accessed. When using a single virtual server for OWA, Outlook Anywhere, ActiveSync, and Autodiscover, you need to use an iRule to separate the traffic that supports cookie persistence (Outlook Web App and ActiveSync) from that which does not (Outlook Anywhere) and assign appropriate persistence methods. This example creates a persistence iRule that uses correct persistence methods for each access type. This iRule assumes the use of separate pools for the services as configured by the template.

#### Critical

*You must change the pool names in the following iRules (shown in red) to match the names of the pools in your configuration*

*Because of the length of this iRule, you can use the following text file to make the copy paste operation easier: <http://www.f5.com/solution-center/deployment-guides/files/exchange-persist.zip>.*

*However, **if you download the zip file, you must still modify the iRule to match the name of the pools in your configuration.***

```

1  ## iRule to select pool and persistence method when all Exchange Client
2  ## Access HTTP-based services are accessed through the same BIG-IP virtual
3  ## server. This iRule will use an HTTP header inserted by a BIG-IP Edge
4  ## Gateway for persistence (if that header is present); otherwise it will
5  ## set persistence according to traditional methods.
6
7  ## CHANGE ALL POOL NAMES TO MATCH THOSE IN YOUR ENVIRONMENT.
8
9  when HTTP_REQUEST {
10
11     ## Offline Address Book and Autodiscover do not require persistence.
12
13     switch -glob -- [string tolower [HTTP::path]] {
14
15         "/microsoft-server-activesync" {
16             ## ActiveSync.
17             if { [HTTP::header exists "APM_session"] } {
18                 persist uie [HTTP::header "APM_session"] 7200
19             } elseif { [HTTP::header exists "Authorization"] } {
20                 persist uie [HTTP::header "Authorization"] 7200
21             } else {
22                 persist source_addr
23             }
24             pool as_pool_name
25             COMPRESS::disable
26             CACHE::disable
27             return
28         }
29
30         "/owa*" {
31             ## Outlook Web Access
32             if { [HTTP::header exists "APM_session"] } {
33                 persist uie [HTTP::header "APM_session"] 7200
34             } else {
35                 persist cookie insert
36             }
37             pool owa_pool_name
38             return
39         }
40
41         "/ecp*" {
42             ## Exchange Control Panel.
43             if { [HTTP::header exists "APM_session"] } {
44                 persist uie [HTTP::header "APM_session"] 7200
45             } else {
46                 persist cookie insert
47             }
48             pool owa_pool_name
49             return
50         }
51
52         "/ews*" {
53             ## Exchange Web Services.
54             if { [HTTP::header exists "APM_session"] } {
55                 persist uie [HTTP::header "APM_session"] 7200
56             } else {
57                 persist source_addr
58             }
59             pool oa_pool_name
60             COMPRESS::disable
61             CACHE::disable
62             return
63         }
64     }
65 }

```

➤ **Critical** This iRule continues on the following page.

⇒ **Critical** This iRule is a continuation of the iRule from the previous page.

```

64     "/oab*" {
65         ## Offline Address Book.
66         pool oa_pool_name
67         return
68     }
69
70     "/rpc/rpcproxy.dll" {
71         ## Outlook Anywhere.
72         if { [HTTP::header exists "APM_session"] } {
73             persist uie [HTTP::header "APM_session"] 7200
74         } elseif { [string tolower [HTTP::header "Authorization"]] starts_with "basic" } {
75             persist uie [HTTP::header "Authorization"] 7200
76         } else {
77             persist source_addr
78         }
79
80         pool oa_pool_name
81         COMPRESS::disable
82         CACHE::disable
83         return
84     }
85
86     "/autodiscover*" {
87         ## Autodiscover.
88         pool ad_pool_name
89         return
90     }
91
92     default {
93         ## This final section takes all traffic that has not otherwise
94         ## been accounted for and sends it to the pool for Outlook Web App
95         if { [HTTP::header exists "APM_session"] } {
96             persist uie [HTTP::header "APM_session"] 7200
97         } else {
98             persist source_addr
99         }
100        pool owa_pool_name
101    }
102 }
103 }
104
105 when HTTP_RESPONSE {
106     if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate" } {
107         ONECONNECT::reuse disable
108         ONECONNECT::detach disable
109         ## this command disables NTLM conn pool for connections where OneConnect has been disabled
110         NTLM::disable
111     }
112     ## this command rechunks encoded responses
113     if {[HTTP::header exists "Transfer-Encoding"]} {
114         HTTP::payload rechunk
115     }
116 }

```

This completes the iRule.

### Outlook Anywhere persistence iRule if using separate pools AND virtual servers

This iRule is necessary because the Microsoft Outlook client does not support HTTP cookies, so the BIG-IP LTM persists based on other HTTP header information. In some cases you may be able to use other persistence methods such as Source Address Affinity, which bases persistence on the IP address of the client. However, because proxy servers or NAT (network address translation) devices may aggregate clients behind a single IP address, such methods are not always effective. To ensure reliable persistence, we recommend using the following iRule and associated persistence profile.

```

1  when HTTP_REQUEST {
2      switch -glob -- [string tolower [HTTP::path]] {
3          "/ews*" {
4              ## Exchange Web Services.
5              if { [HTTP::header exists "APM_session"] } {
6                  persist uie [HTTP::header "APM_session"] 7200
7              } else {
8                  persist source_addr
9              }
10         }
11
12         "/rpc/rpcproxy.dll" {
13             ## Outlook Anywhere.
14             if { [HTTP::header exists "APM_session"] } {
15                 persist uie [HTTP::header "APM_session"] 7200
16             } elseif { [string tolower [HTTP::header "Authorization"]] starts_with "basic" } {
17                 persist uie [HTTP::header "Authorization"] 7200
18             } else {
19                 persist source_addr
20             }
21         }
22     }
23 }
24
25 when HTTP_RESPONSE {
26     if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate" } {
27         ONECONNECT::reuse disable
28         ONECONNECT::detach disable
29         ## disables NTLM conn pool for connections where OneConnect has been disabled
30         NTLM::disable
31     }
32     ## this command rechunks encoded responses
33     if {[HTTP::header exists "Transfer-Encoding"]} {
34         HTTP::payload rechunk
35     }
36 }

```

### Advanced health monitor configuration

If you choose to configure advanced monitors, the BIG-IP performs logins to most of the Client Access services (all except RPC/MAPI) and checks for valid content in the response. Because these monitors attempt to access a specific mailbox, they can more accurately determine the actual health of Client Access services. However, account maintenance and Mailbox status must become a part of your monitoring strategy.

### Important note about BIG-IP health monitors that use Exchange server accounts

The monitors described in this section require a valid Exchange server account and associated mailbox specifically for monitoring purposes. The accounts used for authentication must be associated with a valid mailbox. If authentication should fail for any reason, for instance, the account is locked, the Mailbox server associated with that account is down for maintenance, or the account password is changed, the monitors will mark **all** Client Access servers down for the relevant service (Autodiscover, ActiveSync, or Outlook Anywhere). Maintenance of the accounts and associated mailboxes thus becomes an integral part of your health status checks.

If you choose to use this method, we recommend using at least two separate instances of the monitor, with Mailboxes located on different servers. You should then configure the pool to only mark members down if all monitors fail.

You should create accounts (and associated mailboxes) for monitoring that are not accessed by actual users and that do not have privileged access anywhere else in your network. Because you have to store the user name and password in plain text on your BIG-IPs, make sure the credentials are not used elsewhere in your organization for anything other than monitoring.

We strongly recommend creating a mailbox account(s) specifically for use in the monitor(s).

### Outlook Web App advanced monitor

To create the advance monitor for Outlook Web App, use the following table

Monitor Field	Non-default settings/Notes
<b>Name</b>	Give the monitor a unique name
<b>Type</b>	<b>HTTP</b> (use <b>HTTPS</b> parent if configuring SSL Bridging)
<b>Interval</b>	<b>30</b> (recommended)
<b>Timeout</b>	<b>90</b> (recommended)
<b>Send String<sup>1</sup></b>	If using the default <b>forms-based</b> authentication for OWA <sup>1</sup> <b>GET /owa/auth/logon.aspx?url=https://mail.example.com/owa/&amp;reason=0 HTTP/1.1\r\nUser-Agent: Mozilla/4.0\r\nHost: mail.example.com\r\n\r\n</b> If using <b>Basic</b> or <b>Basic and Windows Integrated Authentication</b> for OWA <b>GET /owa/\r\n</b>
<b>Receive String<sup>2</sup></b>	<b>OutlookSession=</b>
<b>User Name</b>	Type an appropriate user name
<b>Password</b>	Type the associated password
We recommend using this table to create a second monitor, using a second mailbox account.	

<sup>1</sup> Replace red text with your FQDN. This string must be entered as a single line

<sup>2</sup> This response string is part of a Cookie header that OWA returns. Although you may elect to use another string on the page, it must be on the first 5,120 bytes of the received data (including headers and payload). Strings found near the end of the HTTP response from OWA will not be properly detected. See <http://support.f5.com/kb/en-us/solutions/public/3000/400/sol3451.html> for more details.

### Outlook Anywhere advanced monitor

To create the advance monitor for Outlook Anywhere, use the following table

Monitor Field	Non-default settings/Notes
<b>Name</b>	Give the monitor a unique name
<b>Type</b>	<b>HTTP</b> (use <b>HTTPS</b> parent if configuring SSL Bridging)
<b>Interval</b>	<b>30</b> (recommended)
<b>Timeout</b>	<b>90</b> (recommended)
<b>Send String<sup>1</sup></b>	<b>RPC_IN_DATA /rpc/rpcproxy.dll?mail.example.com:6001 HTTP/1.1\r\nUser-Agent: MSRPC\r\nHost: mail.example.com\r\n</b>
<b>Receive String<sup>2</sup></b>	<b>200 Success</b>
<b>User Name</b>	Type an appropriate user name
<b>Password</b>	Type the associated password
We recommend using this table to create a second monitor, using a second mailbox account.	

<sup>1</sup> Replace red text with your FQDN. This string must be entered as a single line. You must only include a single \r\n at the end of the string.

### ActiveSync advanced monitor

To create the advance monitor for ActiveSync, use the following table

Monitor Field	Non-default settings/Notes
<b>Name</b>	Give the monitor a unique name
<b>Type</b>	<b>HTTP</b> (use <b>HTTPS</b> parent if configuring SSL Bridging)
<b>Interval</b>	<b>30</b> (recommended)
<b>Timeout</b>	<b>90</b> (recommended)
<b>Send String<sup>1</sup></b>	<b>GET /Microsoft-Server-ActiveSync/ HTTP/1.1\r\nHost: mail.example.com\r\n</b>
<b>Receive String<sup>2</sup></b>	<b>supported</b>
<b>User Name</b>	Type an appropriate user name
<b>Password</b>	Type the associated password

We recommend using this table to create a second monitor, using a second mailbox account.

<sup>1</sup> Replace red text with your FQDN. This string must be entered as a single line. You must only include a single \r\n at the end of the string.

### Autodiscover advanced EAV monitor

The HTTP monitor for Autodiscover checks for the availability of the web home page. For a more sophisticated health check, it is possible to simulate user login to an actual email inbox using an EAV, or external monitor, instead. The monitor described in this section require a valid Exchange server account and associated mailbox specifically for monitoring purposes.

#### Important



*If you are using a redundant BIG-IP system, you need to make sure any modifications to the script EAVs are manually copied between BIG-IP LTMs, and given the required permissions when configuration is synced.*

First you must download and install the monitor on each BIG-IP system, create the external monitor manually that calls the script, then add the monitor to the load balancing pool.

#### To download and install the script

1. Download the appropriate script:

**If the BIG-IP system is configured for SSL Offload:**

[www.f5.com/solution-center/deployment-guides/files/autodiscover-monitor.zip](http://www.f5.com/solution-center/deployment-guides/files/autodiscover-monitor.zip)

**If the BIG-IP system is configured for SSL Bridging (SSL re-encryption):**

[www.f5.com/solution-center/deployment-guides/files/autodiscover-monitor-ssl-bridging.zip](http://www.f5.com/solution-center/deployment-guides/files/autodiscover-monitor-ssl-bridging.zip)

2. Extract the file to a location accessible by the BIG-IP system.
3. From the Main tab of the BIG-IP Configuration utility, expand **System**, and then click **File Management**.
4. On the Menu bar, click **External Monitor Program File List**.
5. Click the **Import** button.
6. In the **File Name** row, click **Browse**, and then locate the **autodiscover-monitor.sh** (SSL offload) or **autodiscover-monitor-ssl-bridging.sh** (SSL bridging) file.

7. In the **Name** box, type a name for the file. In our example, we type **Autodiscover-monitor**.
8. Click the **Import** button.

The next task is to create the EAV monitor on the BIG-IP system that references the script.

**To create the EAV health monitor that calls the script**

Use the guidance in the following table to create a new external monitor. The table contains all of the non-default settings required for this monitor. For more information on external monitors, or for instructions on configuring the monitor, see the online help or the product documentation.

To start the monitor creation, from the BIG-IP Configuration utility Main tab, expand **Local Traffic**, click **Monitors**, and then click the **Create** button.


Monitor Field	Description/Notes	
<b>Name</b>	User choice	
<b>Type</b>	<b>External</b> (the <i>Import Settings</i> field automatically selects External as well)	
<b>Interval</b>	User choice, but we recommend <b>60</b>	
<b>Timeout</b>	User choice, but we recommend <b>181</b>	
<b>External Program</b>	From the list, select the monitor you imported. In our example, it is <b>Autodiscover-monitor</b> .	
<b>Variables</b>	<i>Name</i>	<i>Value</i>
	<b>USER</b>	The account name associated with a mailbox.
	<b>PASSWORD</b>	The password for the account
	<b>DOMAIN</b>	The Windows domain for the account
	<b>EMAIL</b>	The email address for the user mailbox (such as j.smith@example.com)

Click the **Finished** button. This completes the EAV monitor configuration.

## BIG-IP Edge Gateway and APM manual configuration


This section covers the following scenarios for Edge Gateway and APM:

1. An Edge Gateway deployment on a separate BIG-IP than that providing your Exchange 2010 traffic management. There are two options in this scenario:
  - A. SSL (HTTPS, port 443) connections will be terminated at the Edge Gateway and forwarded to the BIG-IP LTM and then to your Exchange Client Access servers on HTTP port 80.




**Figure 1:** Edge Gateway with SSL Offload configuration example

- B. Both the BIG-IP Edge Gateway and the BIG-IP LTM will perform SSL Bridging; they will decrypt SSL traffic in order to process it, then re-encrypt the traffic before placing it back on the network.




**Figure 2:** Edge Gateway with SSL Bridging configuration example

2. A single BIG-IP configured with both APM and LTM modules. There are two options in this scenario:
  - A. The BIG-IP will terminate SSL connections and forward traffic to your Exchange Client Access servers on HTTP port 80.



**Figure 3:** Edge Gateway with SSL Offload configuration example

- B. The BIG-IP will perform SSL bridging; SSL will be decrypted on the BIG-IP but re-encrypted before it is placed back on the network.



**Figure 4:** BIG-IP APM with SSL Bridging configuration example

## BIG-IP Edge Gateway and APM Configuration

No matter which of the scenarios you are deploying, use the following table to create the BIG-IP Edge Gateway or APM configuration (scenario-specific configuration begins after this section). Create the objects in the order they appear in the table. The tables in this section provide guidance on configuring the individual BIG-IP objects. For specific instructions on configuring individual objects, see the online help or product documentation.

BIG-IP Object	Non-default settings/Notes	
<b>AAA Server</b> (Main tab-->Access Policy-->AAA Servers)	<b>Name</b>	Type a unique name. We use <b>exchange-aaa-server</b> .
	<b>Type</b>	<b>Active Directory</b>
	<b>Domain Controller</b>	Type the IP address or FQDN name of an Active Directory Domain Controller
	<b>Domain Name</b>	Type the Active Directory domain name
	<b>Admin Name<sup>1</sup></b>	Type the AD user name with administrative permissions (optional)
	<b>Admin Password<sup>1</sup></b>	Type the associated password (optional). Type it again in the Verify Password box
<b>SSO Configuration</b> (Main tab-->Access Policy-->SSO Configurations)	<b>Forms based SSO Configuration</b>	
	<b>Name</b>	Type a unique name. We use <b>exchange-forms-sso</b> .
	<b>SSO Method</b>	<b>Form Based</b>
	<b>Form Method</b>	<b>POST</b>
	<b>Form Action</b>	<b>/owa/auth/owaauth.dll</b>
	<b>Form Parameter for User Name</b>	<b>username</b>
	<b>Form Parameter for Password</b>	<b>password</b>
	<b>Start URI</b>	<b>/owa/auth/logon.aspx?url=https://owa.example.com/owa/&amp;reason=0</b> (replace red text with your FQDN)
	<b>Hidden Form Parameters/Values</b>	<b>destination https://owa.example.com/owa/</b> (replace with your FQDN) <b>flags 0</b> <b>forcedownlevel 0</b> <b>isUtf8 1</b> <b>trusted 0</b> (each entry on a separate line)
	<b>NTLM SSO Configuration</b>	
	<b>Name</b>	Type a unique name. We use <b>exchange-ntlm-sso</b> .
	<b>SSO Method</b>	<b>NTLMv1</b>
	<b>Username Conversion</b>	<b>Enable</b>
	<b>NTLM Domain</b>	The NTLM domain name where the user accounts are located
<b>Access Profile</b> (Main tab-->Access Policy-->Access Profiles)	<b>Name</b>	Type a unique name. We use <b>exchange-access</b> .
	<b>SSO Configuration</b>	Select name of NTLM SSO configuration you created above
<b>Access Policy</b> (See procedure below)	<b>Edit</b>	Edit the Access Profile you just created using the Visual Policy Editor Continue now with configuring the Access policy below.

<sup>1</sup> Optional. The Admin Name and Password are only required if anonymous binding to Active Directory is not allowed in your environment.

### Configuring the Access Policy

After creating the objects in the table above, use the following procedure to edit the Access Policy on the Edge Gateway or BIG-IP APM using the Visual Policy Editor (VPE). The Policy shown in the following procedure is just an example, you can use this Access Policy or create one of your own.


#### To configure the Access Policy

1. On the Main tab, expand **Access Policy**, and click **Access Profiles**.
2. Locate the Access Profile you created using the table above, and then, in the Access Policy column, click **Edit**. The VPE opens in a new window.
3. Click the **+** symbol between **Start** and **Deny**. A box opens with options for different actions.

**Note:**

Some of the choices and settings of your Edge Gateway or APM configuration will vary depending on your intended Outlook connectivity method (Outlook Anywhere or MAPI), and your DNS topology (whether you have a single DNS namespace, or separate namespaces for internal and external users, aka 'split-horizon' DNS).

- A. Click the **Logon Page** option button, and then the **Add Item** button at the bottom.
- B. From the **Split domain from full Username** list, select **Yes**.
- C. Configure the rest of the Logon Page properties as applicable, and then click **Save**.
4. Click the **+** symbol between **Logon Page** and **Deny**.
  - A. In the Authentication section, click the **AD Auth** option button, and click **Add Item**.
  - B. In the **Active Directory** properties box, from the **Server** list, select the AAA server you created using the table above. The rest of the settings are optional. Click **Save**.
5. On the Successful path between **AD Auth** and **Deny**, click the **+** symbol.
  - A. Click the **SSO Credential Mapping** option button, and then click **Add Item**.
  - B. Configure the Properties as applicable for your configuration. In our example, we leave the settings at the defaults. Click the **Save** button.
6. On the fallback path between **SSO Credential Mapping** and Deny, click the **Deny** box. Click the **Allow** option button, and then click **Save**. When you are finished, the VPE should look like the image below.
7. Click the yellow **Apply Access Policy** link in the upper left part of the window. You must apply an access policy before it takes effect.



**Creating the iRule that chooses the SSO Configuration**

The next task is to create an iRule that selects the appropriate SSO Configuration to support forms-based authentication of Outlook Web App.

**To create the iRule**

1. On the Main tab, expand **Local Traffic**, click **iRules**, and then click the **Create** button.
2. In the **Name** box, give the iRule a unique name. We use **select\_SSO\_irule**.
3. In the **Definition** section, copy and paste the following iRule, omitting the line numbers. If you used a different name for your forms-based SSO Configuration when creating it based on the table above, use that name in line 3

```

1  when RULE_INIT {
2      ##replace edge_forms_sso here with your forms-based SSO name
3      set static::OWA_FORM_BASE_SSO_CFG_NAME      "exchange_forms_sso"
4  }
5  when ACCESS_ACL_ALLOWED {
6      set req_uri [HTTP::uri]
7      #selects the forms-based SSO when Outlook Web Access is detected
8      if { $req_uri contains "/owa/&reason=0" } {
9          WEBSO::select $static::OWA_FORM_BASE_SSO_CFG_NAME
10     }
11     unset req_uri
12 }

```

4. Click the **Finished** button.

### Configuration table for scenario 1: Edge Gateway sending traffic to a remote BIG-IP LTM

If you are using the BIG-IP Edge Gateway for scenario 1 with either SSL offload or SSL Bridging, use the following table to configure the Edge Gateway. There are additional procedures immediately following this table.

BIG-IP LTM Object	Non-default settings/Notes		
<b>Health Monitor</b> (Main tab-->Local Traffic -->Monitors)	<b>Type</b>	<b>TCP</b>	
	<b>Interval</b>	<b>30</b> (recommended)	
	<b>Timeout</b>	<b>91</b> (recommended)	
<b>Pool</b> (Main tab-->Local Traffic -->Pools)	<b>Health Monitor</b>	Select the monitor you created above	
	<b>Load Balancing Method</b>	<b>Round Robin</b>	
	<b>Address</b>	Type the IP Address of remote BIG-IP LTM virtual server to which this Edge Gateway will forward traffic	
	<b>Service Port</b>	<b>80</b> if offloading SSL <b>443</b> if re-encrypting for SSL Bridging	
<b>iRule</b>	See <i>Creating the iRule on the Edge Gateway on page 65</i> . You must also have created the <i>Append iRule on page 54</i> .		
<b>Profiles</b> (Main tab-->Local Traffic -->Profiles)	<b>HTTP</b> (Profiles-->Services)	Parent Profile	<b>http</b>
		Parent Profile	<b>wan-optimized-compression</b> application/vnd.ms-publisher application/(xls excel msexcel ms-excel x-excel x-xls xmsexcel x-ms-excel vnd.excel vnd.msexcel vnd.ms-excel) application/(word doc msword winword ms-word x-word x-msword vnd.word vnd.msword vnd.ms-word) application/(xml x-javascript javascript x-ecmascript ecmascript) application/(powerpoint msppowerpoint ms-powerpoint x-powerpoint x-mspowerpoint vnd.powerpoint vnd.mspowerpoint  vnd.ms-powerpoint vnd.ms-pps) application/(mpp msproject x-msproject x-ms-project vnd.ms-project) application/(visio x-visio vnd.visio vsd x-vsd x-vsd) application/(pdf x-pdf acrobat vnd.pdf)
	<b>HTTP Compression</b> (Profiles-->Services)	Content List-->Include List <sup>2</sup> (Add each entry to the Content Type box and click Include)	
	<b>Web Acceleration</b> (Profiles-->Services)	Parent Profile	<b>optimized-caching</b>
	<b>TCP WAN</b> (Profiles-->Protocol)	Parent Profile	<b>tcp-wan-optimized</b>
	<b>TCP LAN</b> (Profiles-->Protocol)	Parent Profile	<b>tcp-lan-optimized</b>
	<b>OneConnect</b> (Profiles-->Other)	Parent Profile	<b>oneconnect</b>
	<b>Client SSL</b> (Profiles-->SSL)	Parent Profile Certificate and Key	<b>clientssl</b> Select your Certificate and key
	<b>Server SSL</b> <i>(for SSL Bridging only)</i> (Profiles-->SSL)	Parent Profile Certificate and Key	If the remote BIG-IP LTM receiving this traffic is using a self-signed or default certificate for decryption, select <b>serverssl-insecure-compatible</b> If the remote LTM is using a certificate signed by a Certificate Authority, select <b>serverssl</b> Select your Certificate and key
	<b>Edge Gateway Exchange virtual server</b> (Main tab-->Local Traffic-->Virtual Servers)	<b>Name</b>	Type a unique name. We use <b>edge-exchange-virtual</b> .
	<b>Destination Address</b>	The IP address clients use to access Microsoft Exchange. Your Exchange FQDN resolves to this IP address.	
	<b>Service Port</b>	<b>443</b>	
	<b>OneConnect profile</b>	Select the OneConnect profile you created above.	
	<b>HTTP Profile</b>	Select the HTTP profile you created above	
	<b>HTTP Compression Profile</b>	Select the HTTP Compression profile you created above.	
	<b>Web Acceleration Profile</b>	Select the Web Acceleration profile you created above	

This table continues on the following page

BIG-IP LTM Object	Non-default settings/Notes
<b>Edge Gateway Exchange virtual server</b> (Main tab-->Local Traffic-->Virtual Servers)	<b>SSL Profile (Client)</b> Select the Client SSL profile you created above.
	<b>SSL Profile (Server)</b> Select the Server SSL profile you created above (only for Scenario 2, SSL Bridging).
	<b>Access Profile</b> Select the Access Profile you created above
	<b>iRules</b> Enable the <b>Append</b> iRule you created on page 54. Enable the built-in <b>_sys_APM_ExchangeSupport_OA_BasicAuth</b> iRule. This rule is necessary whether deploying Outlook Anywhere or not. Enable the iRule that chooses the SSO configuration you created ( <b>select_SSO_irule</b> in our example) Enable the APM session ID irule you created ( <b>edge-gateway-irule</b> in our example)
	<b>Default Pool</b> Select the Pool you created above

### Creating the iRule on the Edge Gateway

The first task is to create the iRule on the BIG-IP LTM for Edge Gateway. The first iRule is necessary for all deployments with Edge Gateway.

#### To create the iRule to persist connections based on APM session ID on the Edge Gateway

1. On the Main tab, expand **Local Traffic**, click **iRules**, and then click the **Create** button.
2. In the **Name** box, give the iRule a unique name. We use **edge-gateway-irule**.
3. In the **Definition** section, copy and paste the following iRule, omitting the line numbers.

```

1  when ACCESS_ACL_ALLOWED {
2      set sessionid [ACCESS::session data get "session.user.sessionid"]
3      HTTP::header insert APM_session $sessionid
4  }
5  when HTTP_RESPONSE {
6      if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate"} {
7          ONECONNECT::reuse disable
8          ONECONNECT::detach disable
9          ## disables NTLM conn pool for connections where OneConnect has been disabled
10         NTLM::disable
11     }
12     ## this command rechunks encoded responses
13     if {[HTTP::header exists "Transfer-Encoding"]} {
14         HTTP::payload rechunk
15     }
16 }
    
```

4. Click the **Finished** button.

#### BIG-IP LTM iRule if all traffic goes through the Edge Gateway

If all of your Exchange 2010 traffic goes through the Edge Gateway, and you do not have internal users who go directly to the BIG-IP LTM, you must modify the persistence iRule on the remote BIG-IP LTM to use the following iRule (and remove the existing persistence iRule).

**Important**



*This iRule is only necessary if all traffic is going through the Edge Gateway. If you have internal users who go directly to the BIG-IP LTM, **do not** use this iRule.*

#### To create the persistence iRule if all traffic goes through the Edge Gateway to the LTM

1. On the Main tab, expand **Local Traffic**, click **iRules**, and then click the **Create** button to create a new iRule if all traffic goes through the Edge Gateway.
2. In the **Name** box, type a unique name. In our example, we type **edge-gateway-persist**.
3. In the **Definition** section, copy and paste the following iRule, omitting the line numbers.

```

1  when HTTP_REQUEST {
2
3  ## CHANGE ALL POOL NAMES TO MATCH THOSE IN YOUR ENVIRONMENT.
4  ## OAB and Autodiscover do not require persistence.
5
6      switch -glob -- [string tolower [HTTP::path]] {
7
8          "/microsoft-server-activesync" {
9              pool my_Exchange_2010__single_as_pool
10             COMPRESS::disable
11             CACHE::disable
12             persist uie [HTTP::header "APM_session"] 7200
13             return
14         }
15
16         "/ews*" {
17             pool my_Exchange_2010__single_owa_pool
18             COMPRESS::disable
19             CACHE::disable
20             persist uie [HTTP::header "APM_session"] 7200
21             return
22         }
23
24         "/ecp*" {
25             pool my_Exchange_2010__single_owa_pool
26             persist uie [HTTP::header "APM_session"] 7200
27             return
28         }
29
30         "/oab*" {
31             pool my_Exchange_2010__single_owa_pool
32             return
33         }
34
35         "/rpc/rpcproxy.dll" {
36             pool my_Exchange_2010__single_oa_pool
37             COMPRESS::disable
38             CACHE::disable
39             persist uie [HTTP::header "APM_session"] 7200
40             return
41         }
42
43         "/autodiscover*" {
44             pool my_Exchange_2010__single_ad_pool
45             return
46         }
47
48         default {
49             ## This final section takes all traffic that has not otherwise
50             ## been accounted for and sends it to the pool for Outlook Web
51             ## App
52             pool my_Exchange_2010__single_owa_pool
53             persist uie [HTTP::header "APM_session"] 7200
54         }
55     }
56 }
57 when HTTP_RESPONSE {
58     if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate" } {
59         ONECONNECT::reuse disable
60         ONECONNECT::detach disable
61         ## disables NTLM conn pool for connections where OneConnect has been disabled
62         NTLM::disable
63     }
64     ## this command rechunks encoded responses
65     if {[HTTP::header exists "Transfer-Encoding"]} {
66         HTTP::payload rechunk
67     }
68 }

```

4. Click the **Finished** button.

#### **Modifying the virtual server to use the new persistence iRule**

If you just created the new persistence iRule on the BIG-IP LTM (*BIG-IP LTM iRule if all traffic goes through the Edge Gateway on page 65*), and have an existing BIG-IP LTM configuration, you must modify the BIG-IP LTM virtual server to use the new persistence iRule and remove any existing persistence iRules.

This completes the Edge Gateway configuration for scenario 1.

#### **Configuration for scenario 2: Single BIG-IP with LTM and APM**

If you are configuring the BIG-IP APM as a module on the same physical BIG-IP device as the LTM configuration, you must modify your BIG-IP LTM configuration to use the following persistence iRule, and remove any existing persistence iRules on the LTM.

#### **Creating the persistence iRule when using BIG-IP APM**

The next task is to create a new persistence iRule on the BIG-IP system for APM.

##### **To create the iRule**

1. On the Main tab, expand **Local Traffic**, click **iRules**, and then click the **Create** button.
2. In the Name box, give the iRule a unique name. We use **apm-persistence-irule**.
3. In the **Definition** section, copy and paste the iRule on the following page, omitting the line numbers.

```

1  ## CHANGE ALL POOL NAMES TO MATCH THOSE IN YOUR ENVIRONMENT.
2  ## OAB and Autodiscover do not require persistence.
3
4  when ACCESS_ACL_ALLOWED {
5      set sessionid [ACCESS::session data get "session.user.sessionid"]
6
7      switch -glob -- [string tolower [HTTP::path]] {
8
9          "/microsoft-server-activesync" {
10             pool my_Exchange_2010__single_as_pool
11             COMPRESS::disable
12             CACHE::disable
13             persist uie $sessionid 7200
14             return
15         }
16         "/ews*" {
17             pool my_Exchange_2010__single_owa_pool
18             COMPRESS::disable
19             CACHE::disable
20             persist uie $sessionid 7200
21             return
22         }
23         "/ecp*" {
24             pool my_Exchange_2010__single_owa_pool
25             persist uie $sessionid 7200
26             return
27         }
28         "/oab*" {
29             pool my_Exchange_2010__single_owa_pool
30             return
31         }
32         "/rpc/rpcproxy.dll" {
33             pool my_Exchange_2010__single_oa_pool
34             COMPRESS::disable
35             CACHE::disable
36             persist uie $sessionid 7200
37             return
38         }
39         "/autodiscover*" {
40             pool my_Exchange_2010__single_ad_pool
41             return
42         }
43
44         default {
45             ## This final section takes all traffic that has not otherwise
46             ## been accounted for and sends it to the pool for Outlook Web
47             ## App
48             pool my_Exchange_2010__single_owa_pool
49             persist uie $sessionid 7200
50         }
51     }
52 }
53 when HTTP_RESPONSE {
54     if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate" } {
55         ONECONNECT::reuse disable
56         ONECONNECT::detach disable
57         ## disables NTLM conn pool for connections where OneConnect has been disabled
58         NTLM::disable
59     }
60     ## this command rechunks encoded responses
61     if {[HTTP::header exists "Transfer-Encoding"]} {
62         HTTP::payload rechunk
63     }
64 }

```

4. Click **Finished**.

### **Modifying the virtual server to use the APM persistence iRule and Access Profile**

The final task is to modify the BIG-IP LTM virtual server(s) to use the new persistence iRule and remove any existing persistence iRules, and add the Access Profile you created on BIG-IP APM.

If you created separate virtual servers, you must add the iRule and Access Profile to all BIG-IP LTM virtual server for the HTTP-based Client Access Services (Outlook Web App, Outlook Anywhere, ActiveSync, and Autodiscover).

## Appendix D: Technical Notes

The following contains additional information that may be helpful when configuring the BIG-IP system for Microsoft Exchange Server 2010.

### Slow Ramp Time

When you configure a Slow Ramp time, BIG-IP will not immediately send a full proportional share of incoming traffic to a pool member that has just come online. Instead, the BIG-IP will increase the proportion of traffic gradually over the time specified. This ensures that a newly-booted or newly-added server is not overwhelmed with incoming traffic, especially when you have selected a Least Connections load-balancing method.

Although advanced monitors that perform logins will prevent any traffic being sent to a Client Access server until at least those functions are enabled, other background services may not be fully ready to service connections. As such, we strongly recommend Slow Ramp even with advanced monitors. If you are not using advanced monitors but have only enabled simple TCP checks or HTTP queries that do not actually check for full client functionality, a Slow Ramp time is essential.

F5 testing has shown that 300 seconds (5 minutes) is generally sufficient to allow a rebooted Exchange 2010 Client Access server to fully start all services and be ready to handle a full load of traffic, but that time is highly dependent on local conditions. You may want to adjust the time period up or down in your environment based on your server capacity and load.

### Subject Alternative Name (SAN) SSL Certificates

This template currently only supports the use of a single DNS name and corresponding certificate and key for all services, or multiple DNS names using a SAN-enabled certificate and key. Support for multiple names, each with separate corresponding certificates and keys, will be in a future release.

An SSL certificate that supports the Subject Alternative Name (SAN) extension allows more than one valid FQDN per certificate, without having to resort to a "wildcard" certificate for a domain. When used in conjunction with Exchange Server 2010, SAN certificates make it simple to combine multiple services into a single virtual server while retaining the flexibility of separate FQDNs. Some examples of using SAN certificates with Exchange 2010 are shown in this [TechNet Article](#).

When you request a SAN certificate from a certification authority, you must define all desired FQDNs in the Subject Alternative Name field; clients will ignore the Common Name in the certificate Subject.

In BIG-IP versions prior to 11.1, the BIG-IP web-based Configuration utility does not display the Subject Alternative Name values of imported certificates, however, the use of SAN certificates is otherwise supported.

### Maximum number of concurrent users: SNAT Pool guidance

If you expect fewer than 6,000 concurrent users per Client Access Server, the iApp configures SNAT Automap. If you expect more than 6,000 users, the iApp configures a SNAT Pool. This section describes how F5 chose 6,000 users as a rule of thumb, and contains additional information if you want to more precisely calculate the number of concurrent users for your SNAT Pool configuration.

The BIG-IP system can create roughly 64,000 connections per SNAT address (ephemeral or source ports used by connections from the BIG-IP range from 1024 to 65,535, or an absolute maximum 64,511 effective concurrent connections). Each user connected to a Client Access server can have about 10 concurrent connections (for example, if a user has Outlook on a PC, a mobile phone,

and Lync running simultaneously). Therefore, you would need a SNAT address for each 6,000 concurrent users you expect. For example, if you have 12,000 users, you need two SNAT pool IP addresses; if you have 15,000 users, you need three addresses. The IP address(es) you specify must not be self IP addresses on this BIG-IP system.

The following table shows the number of expected connections for different clients. To calculate the number of connections, you should determine the approximate number and mix of client types and access methods, using the following table as a guide.

Client	Concurrent Connections (default/typical)
Outlook 2010 (RPC Client Access/MAPI)	Up to 20
Outlook 2010 (Outlook Anywhere)	10
Internet Explorer 7	2
Internet Explorer 8	6
Internet Explorer 9	2
Firefox	6
Chrome	6
Safari	8

### Outlook Client Configuration

Exchange administrators will typically use Autodiscover to configure Outlook clients. If manual configuration is required, the following table provides the recommended settings to match the deployment scenarios described in this guide.

Connection Settings	Default	Your Setting	Notes
Connect to Microsoft Exchange using HTTP	Not selected	Selected	This enables Outlook Anywhere
Use this URL to connect to my Proxy server for Exchange	No default value	FQDN of your Outlook Anywhere virtual server on your Edge Gateway	
Connect using SSL only	Selected	Selected	
On fast networks, connect using HTTP first, then connect using TCP/IP	Not selected	Selected	
On slow networks, connect using HTTP first, then connect using TCP/IP	Selected	Selected	
Proxy authentication settings	NTLM	Basic	

## Creating a new Client Access Array

To create a new Client Access Array, use the Exchange Management Shell to run the following command:

```
New-ClientAccessArray -Name "ArrayName" -FQDN outlook.example.com -Site "SiteName"
```

You must replace *ArrayName* with the name you want for your Client Access Array, replace *outlook.example.com* with the FQDN you have configured in DNS, and replace *SiteName* with the name of your Active Directory site.

You must modify the attributes of any pre-existing mailbox databases to use the new array. Use the Exchange Management Shell to run the following command for each database in your array:

```
Set-MailboxDatabase "MailboxDatabaseName" -RPCClientAccessServer outlook.example.com
```

You must specify the correct mailbox databases for your site, and the correct FQDN for your Client Access Array. You can only configure one Client Access Array (and thus one FQDN and one BIG-IP virtual server) per site.

For complete documentation from Microsoft, see <http://technet.microsoft.com/en-us/library/ee332317.aspx>

## Document Revision History

Version	Description	Date
1.0	New deployment guide for the iApp included in BIG-IP 11.0	N/A
1.1	In the manual configuration appendix (and the downloadable iRule), modified the persistence iRule with the following changes: <ul style="list-style-type: none"> <li>- Changed "/Microsoft-Server-ActiveSync" to "/microsoft-server-activesync" (note case)</li> <li>- Changed "/xml/autodiscover.aspx" to "/autodiscover*"</li> <li>- In the "/rpc/rpcproxy.dll" section, added <b>string tolower</b> to the first else statement.</li> <li>- Changed "MSRPC" to "msrpc"</li> <li>- Changed "*Microsoft Office" to "*microsoft office"</li> </ul>	N/A
1.2	Added the persistence iRule to a ZIP file to avoid any potential white space errors.	N/A
1.3	<ul style="list-style-type: none"> <li>- Added instructions on removing the RPC Client Access Referral Service objects if Static Ports were selected in the template. Updated the manual tables for RPC Client Access.</li> <li>- Updated the SNAT Pool iRule.</li> <li>- Modified the persistence iRule if using a single virtual server (lines 102 and 103 on page 45) to change <b>header</b> to <b>cookie</b>.</li> <li>- Clarified the note about using Internet Explorer 7.0 and 8.0 when configuring the iApp</li> <li>- Added a column for the revision date on this table.</li> </ul>	01/04/2012
1.4	Added an Idle Timeout value of 7200 to the TCP profiles for RPC Client Access, to the required post template configuration and to the manual configuration tables.	01/17/2012
1.5	<ul style="list-style-type: none"> <li>- Modified the optional section on using X-Forwarded-For to log the client IP address in IIS 7 and 7.5 to include installing the Custom Logging service role, and steps for editing the IIS Log Definition to include the X-Forwarded-For header.</li> <li>- Moved the X-Forwarded-For section to Appendix C, <i>page 67</i>.</li> <li>- Added a note to the manual configuration stating that if you download the persistence iRule, you must still modify the pool names in the file you downloaded.</li> <li>- Clarified TCP Request Queuing description in multiple sections.</li> </ul>	03/13/2012
<b>2.0</b>	Completely updated deployment guide for the downloadable iApp for 11.0, 11.0.1, and 11.1. All users must use the downloadable iApp for Exchange 2010 deployments.	04/06/2012
2.1	<ul style="list-style-type: none"> <li>- <i>Manual configuration only:</i> Corrected the persistence iRule for the single virtual server deployment in the manual configuration table. The previous version had variables in the lines prior to the pool names.</li> <li>- Updated the link to Ask F5 to point directly to the Exchange Server 2010 iApp solution (SOL13497) rather than the Index solution (SOL13422).</li> <li>- Clarified the BIG-IP APM configuration is applicable to all Exchange HTTP-based Client Access services, even though the iApp only asks for the FQDN for Outlook Web App.</li> </ul>	04/16/2012
2.2	<i>Manual configuration only:</i> Modified the HTTP Compression profile section of the single virtual server for Exchange HTTP-based services table and OWA configuration table (for separate virtual servers) to remove the requirement to check the <b>Keep Accept Encoding</b> box. You should not enable Keep Accept Encoding. The iApp configures this option correctly.	04/19/2012
2.3	<ul style="list-style-type: none"> <li>- Added a Troubleshooting section that describes what to do if users are unable to connect to Exchange Client Access Servers after deploying the iApp, and if an Edge Gateway implementation is not redirecting users to the OWA URI.</li> <li>- Modified the manual configuration tables with the fixes described in the Troubleshooting section.</li> </ul>	05/03/2012
2.4	<ul style="list-style-type: none"> <li>- Replaced Troubleshooting section with <i>Modifying the iApp template configuration on page 35</i>. This section now contains required and optional change to make to the configuration produced by the iApp.</li> <li>- Updated the Persistence iRules.</li> <li>- Updated the Autodiscover EAV script file, and added a new script file if configuring SSL Bridging.</li> </ul>	05/18/2012

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

