Deployment Guide

Document Version: 4.9.2

Deploying the BIG-IP System v10 with Microsoft Exchange Server 2010

Welcome to the F5 and Microsoft® Exchange® 2010 deployment guide. This document contains guidance on configuring the BIG-IP system version 10.2.1 and later in the v10 branch for Microsoft Exchange 2010, including SP1 and SP2. If you are using the BIG-IP system version 11 or later, see *http://www.f5.com/pdf/deployment-guides/microsoft-exchange2010-iapp-dg.pdf*.

For more information on the F5 devices included in this guide, see *http://www.f5.com/products/*.

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*.

For more information on Microsoft Exchange Server 2010, see *http://www.microsoft.com/exchange/2010/en/us/default.aspx*

Products and versions tested

Product	Version
BIG-IP LTM and Virtual Edition	10.2.1, 10.2.2, 10.2.4
Microsoft Exchange Server	2010 and 2010 SP1, SP2, SP3

➲ **Important:** *Make sure you are using the most recent version of this deployment guide, available at http://www.f5.com/pdf/deployment-guides/f5-exchange-2010-dg.pdf.*

See *Deployment Guide Revision History on page 82* for a description of the document revisions.

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

Microsoft Partner Network

Contents

Chapter 1 Deploying the BIG-IP System with Exchange 2010 Client Access Servers

Configuring F5 devices with the Microsoft Exchange Server 2010 Client Access Role using the Application template

This chapter provides guidance for using the BIG-IP Application Template found in BIG-IP version 10.2.1 and later in the v10 branch to configure the Client Access server role of Microsoft Exchange Server 2010. By using the 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.

For more information on the Client Access Server role, see *technet.microsoft.com/en-us/library/bb124915%28EXCHG.140%29.aspx*

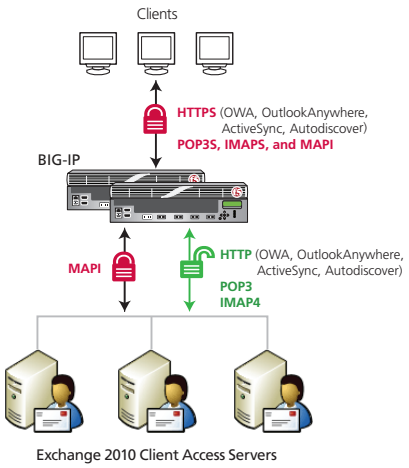

Prerequisites and configuration notes

The following are prerequisites and configuration notes for the Client Access Role:

- 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 full 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.
- We recommend saving the existing BIG-IP configuration before you begin this Deployment Guide. For specific instructions, refer to the manual appropriate for your BIG-IP version, available on Ask F5 (*http://support.f5.com/*)
- 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.
- You must be using BIG-IP version 10.2.1. ***We strongly recommend v10.2.4 or higher.***

Configuration example

In the simplified logical configuration diagram on the left, we show connectivity options for several types of clients to the same Exchange Server 2010 Client Access servers. Users connect directly to the LTM systems via secure connections (HTTPS, MAPI, POP3S, or IMAPS, depending on choice of web browser or email client).

In our example, we show the BIG-IP LTM offloading all SSL processing from the Exchange Client Access servers, while secure MAPI connections (which do not use SSL) are forwarded without being decrypted. Your implementation may be different from the one shown.

Critical →

The instructions in this chapter are valid only for versions 10.2.1 and later in the v10 branch. If you are running an earlier version of the BIG-IP system you should upgrade before running the template, or configure the required objects manually according to the configuration tables in Appendix A: Manual Configuration Tables on page 50.

Using the Application Template for Client Access Servers

The Application Template greatly simplifies configuring the BIG-IP system for Microsoft Exchange 2010 Client Access server roles. The following is based on the Application Template available in BIG-IP version 10.2.1, which has been updated with functionality not found in earlier versions.

Before beginning the Application template, there are some decisions you must make.

- **Which Client Access services are you planning to use?**
  The Exchange 2010 Client Access role contains a number of services. Before starting the Application Template, you must know which of the following services you are using in your Exchange 2010 environment:
  - Outlook Web App (this configuration is also used by the ECP, EWS and OAB services)
  - Outlook Anywhere (RPC over HTTP)
  - Autodiscover
  - POP3
  - ActiveSync
  - RPC Client Access (MAPI)
  - IMAP4

- **Will each service have a dedicated IP address (BIG-IP virtual server) or will all HTTP-based services share an IP address?**
  There are two ways you can configure the BIG-IP system for the Outlook Web App, Outlook Anywhere, ActiveSync, and Autodiscover services:

  - *Separate IP addresses/virtual servers for each HTTP service*
    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 those services are each 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.

  - *One IP address/virtual server for all HTTP services*
    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. See *Subject Alternative Name (SAN) SSL Certificates on page 79* for further details

    If you are using the BIG-IP Edge Gateway or BIG-IP APM for secure access to Client Access servers, you must use a single IP address. See page 31 for details.

Note → *The single virtual server option is only applicable if you are using the BIG-IP LTM for two or more of the following services: Outlook Web App, Outlook Anywhere, ActiveSync, and Autodiscover. The question in the Template does not appear until you have selected Yes to two or more services.*

- **Are you using the WebAccelerator module for Outlook Web App?**
  The WebAccelerator is an advanced web application delivery solution that provides a series of intelligent technologies designed to overcome problems with browsers, web application platforms and WAN latency issues which impact user performance.
  If you plan to use the WebAccelerator module with Outlook Web App, you must have licensed and provisioned the module before beginning the template. For more information, contact your F5 sales representative.

Important → *After completing the application template, you must perform a few additional required steps, and have the option of adding an EAV (extended application verification) monitor for the Autodiscover service. See Modifying the template configuration on page 22 for details.*

Preparation worksheets

For each section of the Application 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. 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.

➲ **Note:** *Although we show space for seven pool members for each service, you may have more or fewer members in each pool.*

Client Access Services that use HTTP

IP Addresses[1]	Pool Members	FQDN	WAN or LAN	WebAccelerator	SSL Processing [2]
Outlook Web App *(see Configuring the BIG-IP system for Outlook Web App on page 8)*					
Virtual server IP address:	IP addresses of the Client Access Servers that are running OWA: 1: 2: 3: 4: 5: 6: 7:	The fully qualified domain name(s) clients are expected to use (used in the health monitor and Web Accelerator):	Most clients connecting through BIG-IP are coming over a: LAN WAN	Are you using the WebAccelerator module? Yes No	Certificate: Key: :
Outlook Anywhere *(see Configuring the BIG-IP system for Outlook Anywhere on page 10)*					
Virtual server IP address:	IP addresses of the servers that are running Outlook Anywhere: 1: 2: 3: 4: 5: 6: 7:	The fully qualified domain name clients are expected to use	Most clients connecting through BIG-IP are coming over a: LAN WAN	Not Applicable	Certificate: Key:
ActiveSync *(see Configuring the BIG-IP system for ActiveSync on page 12)*					
Virtual server IP address:	IP addresses of the servers that are running ActiveSync: 1: 2: 3: 4: 5: 6: 7:	The fully qualified domain name clients are expected to use	Most clients connecting through BIG-IP are coming over a: LAN WAN	Not Applicable	Certificate: Key:
Autodiscover *(see Configuring the BIG-IP system for Autodiscover on page 14)*					
Virtual server IP address:	IP addresses of the servers that are running Autodiscover: 1: 2: 3: 4: 5: 6: 7:	The fully qualified domain name clients are expected to use	Most clients connecting through BIG-IP are coming over a: LAN WAN	Not Applicable	Certificate: Key:

[1] If you are using one IP address for all the Client Access services that use HTTP (see bullet point on previous page), you do not need to add IP addresses for each row.

[2] ***Important:*** ***Before running the template, you must have imported a certificate and key onto the BIG-IP LTM for each FQDN***. For details, see step 7 on *page 9*

RPC Client Access

IP Addresses	Dynamic or static ports	Pool Members	WAN or LAN
RPC Client Access *(see Configuring the BIG-IP system for RPC Client Access on page 16)*			
Virtual server IP address:	By default, the template uses a dynamic range of ports for RPC Client Access. If you want to use static ports: MAPI port: Address Book port:	IP addresses of the Client Access Servers that are running RPC Client Access: 1: 2: 3: 4: 5: 6: 7:	Most clients connecting through BIG-IP are coming over a: LAN WAN

POP3 and IMAP4

IP Addresses	Pool Members	WAN or LAN	Health Monitor	SSL Offload
POP3 *(see Configuring the BIG-IP system for POP3 on page 18)*				
Virtual server IP address:	IP addresses of the Client Access Servers that are running POP3 : 1: 2: 3: 4: 5: 6: 7:	Most clients connecting through BIG-IP are coming over a: LAN WAN	Optional: POP3 user account can be used for health monitor. Domain name for the account in Active Directory: User name Password	Certificate: Key:
IMAP4 *(see Configuring the BIG-IP system for IMAP4 on page 20)*				
Virtual server IP address:	IP addresses of the Client Access servers that are running IMAP4 : 1: 2: 3: 4: 5: 6: 7:	Most clients connecting through BIG-IP are coming over a: LAN WAN	Optional: IMAP4 user account can be used for health monitor. Domain name for the account in Active Directory: User name Password	Certificate: Key:

Using the template to configure Client Access services

In this section, we provide guidance on configuring the BIG-IP system using the Application Template. To access the Template, log on to the BIG-IP system, expand **Templates and Wizards**, click **Templates**, and then click **Microsoft Exchange 2010**.

Global Client Access questions

The first section of the template asks two questions:

1. **Unique prefix**
   The system attaches this prefix to all of the BIG-IP objects created by the template. You can leave the default or create a prefix specific to your implementation.
2. **Manual routes or secure network address translation**
   If the Client Access servers do not have a route back to the clients through the BIG-IP (typical and the default), i.e. if they do not use the BIG-IP as a gateway to client networks, the BIG-IP uses Secure Network Address Translation (SNAT) Automap to translate the client's source address to an address configured on the BIG-IP. The servers then use this new source address as the destination address.

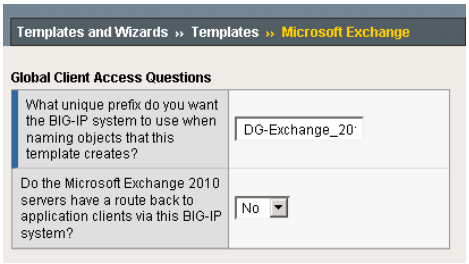

If you indicate that the Client Access servers do have a route back to the clients through the BIG-IP, the BIG-IP does not translate the client's source address; in this case, you must make sure that the BIG-IP is configured as the gateway to the client networks (usually the default gateway) on the Client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

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

If you are configuring your BIG-IP LTM in a "one-armed" configuration with your Client Access servers -- where the BIG-IP virtual server(s) and the Client Access server have IP addresses on the same subnet – you must choose **No**.

➲ **Note:** *For some Exchange 2010 deployments, you may need to use a SNAT Pool instead of SNAT Automap. For more information on SNAT Pools, see Using a SNAT Pool if you expect more than 6,000 users per Client Access server on page 26.*

The following sections contain guidance for each of the Client Access services, in the order they appear in the template.

Configuring the BIG-IP system for Outlook Web App

Outlook Web App (OWA) allows authorized users to securely access their Exchange mailboxes through a browser. The BIG-IP virtual server that you create for OWA is also used for Offline Address Book (OAB), Exchange Control Panel (ECP), and Exchange Web Services (EWS).
By using BIG-IP in front of Outlook Web App servers, you gain the following benefits:

- Terminating HTTPS connections at the BIG-IP LTM reduces CPU and memory load on Outlook Web App servers.
- Terminating HTTPS connections at the BIG-IP simplifies TLS/SSL certificate management.
- The BIG-IP LTM can balance load and ensure high-availability across multiple Outlook Web App servers using a variety of load-balancing methods and priority rules.
- The BIG-IP LTMs 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.

Important →

To configure the Outlook Web App servers to support SSL offloading, you must configure the OWA, ECP, OAB and EWS services according to the documentation at
http://social.technet.microsoft.com/wiki/contents/articles/how-to-configure-ssl-offloading-in-exchange-2010.aspx
We recommend using the scripted method described in that document.

Tip →

*In the Outlook Web App, Outlook Anywhere, ActiveSync, and Autodiscover Questions box, we recommend you select **Yes** for each services you are planning to use. If you select Yes for multiple options, the option for a single or separate IP addresses appears. Choose the option applicable for you configuration. We recommend a single IP address.*

Outlook Web App Virtual Server Questions

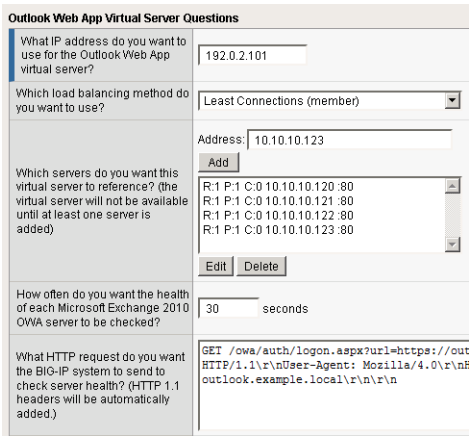

1. **IP address for the virtual server**
   The FQDN for your Outlook Web App service will resolve to this IP address. It must be accessible from external networks, and resolvable via DNS
2. **Load balancing method**
   While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (member)**.
3. **Address**
   Use the IP address for the Client Access servers running Outlook Web App. The template adds the nodes to the appropriate load balancing Pool.
4. **Health Monitor Questions**
   a. *Interval*: Specifies how often the system checks the health of the servers. We recommend the default of 30 seconds.
   b. *HTTP Request*: This is optional but recommended.
   - If you are using the default **forms-based** authentication for OWA, you can configure the template to retrieve the OWA login page and check for a valid response. Replace the default value (GET /) with the following string:

Notes →

- Replace the text in red with the FQDN of your OWA virtual server
- This string must be entered into the text field as one continuous line.

```
GET /owa/auth/logon.aspx?url=https://mail.example.com/owa/&reason=0 HTTP/1.1\r\
nUser-Agent: Mozilla/4.0\r\nHost: mail.example.com\r\n\r\n
```

- If you switched the authentication method to **Basic**, or **Basic and Windows Integrated authentication** (if using Edge Gateway or BIG-IP APM, the auth method must be **Forms**), use `GET /owa/\r\n`. In this scenario, you must modify the monitor configuration after completing the template to provide a valid User Name and Password; the BIG-IP automatically inserts a "\r\n\r\n" after sending the

Authentication header, so unlike the anonymous Forms-based method above, there is only a single "/r/n" at the end of the string. For instructions on modifying the monitor, see *page 27*.

c. *HTTP version*: If you are using the default monitor string, you should leave this at version 1.0. If you are using the custom string above, select **Version 1.1** from the list.

- FQDN: When you select Version 1.1, a new row appears asking for the FQDN for Outlook Web App. Type it here.

d. *Monitor response string*: Optional but recommended. Type the response you expect from the string you entered in step b. If you configured the unique HTTP Request in step b, type a string that is only returned if OWA is functioning. We suggest **OutlookSession=**.
***See note on the lower left for important information on this string**.

Important: *A default installation of Outlook Web App requires SSL on the server and does not have an HTTP (port 80) listener. As noted in the previous section, you must follow Microsoft documentation on configuring Outlook Web App for SSL offloading and creating a port 80 listener on your Client Access servers. If you do not, this monitor will not function properly.*

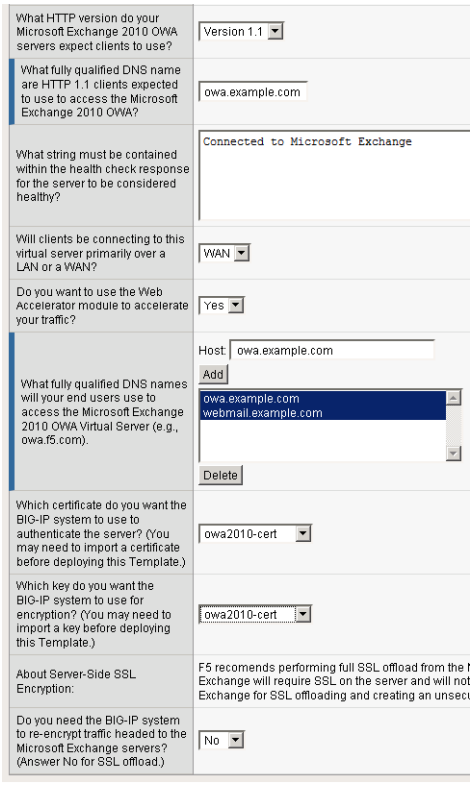

5. **WAN or LAN**
Specify whether most clients are connecting over a WAN or LAN. Because most OWA clients are likely to be coming over the WAN, we recommend selecting WAN (the default).

6. **WebAccelerator**
If you have licensed and provisioned the WebAccelerator module, you have the option of using it for OWA. WebAccelerator provides application acceleration for remote users.

a. When you select Yes, an additional row appears in the template asking for all fully qualified domain names used for Outlook Web App. The BIG-IP system uses these entries for the Requested Hosts field, allowing the WebAccelerator module to accelerate the traffic to these virtual hosts.

7. **SSL Offload**
Before running the template you should have already imported a certificate and key onto the BIG-IP system. While the BIG-IP system does include a self-signed SSL certificate that can be used internally or for testing, we strongly recommend importing a certificate and key issued from a trusted Certificate Authority.
For information on SSL certificates on the BIG-IP system, see the online help or the ***Managing SSL Certificates for Local Traffic*** chapter in the ***Configuration Guide for BIG-IP Local Traffic Manager*** available at *http://support.f5.com/kb/en-us.html*.

» *Certificate*: Choose the certificate for Outlook Web App

» *Key*: Choose the key for Outlook Web App.

8. **SSL re-encryption**
While we recommend offloading SSL from the Client Access servers, some organizations require Exchange traffic is encrypted all the way to the Client Access servers, and not offloaded on the BIG-IP. The BIG-IP system has the ability to re-encrypt the traffic before sending it to the Client Access servers.
To get the greatest benefit from SSL offload, select **No** from the list.
If you require encrypted traffic from the BIG-IP to the Client Access Servers, select **Yes** from the list, and select the appropriate certificate and key.

***Note on monitor response string:**
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.

Configuring the BIG-IP system for Outlook Anywhere

Outlook Anywhere for Exchange 2010 allows you to use Microsoft Outlook clients to connect to your Exchange server over the Internet, using HTTPS to encapsulate RPC (MAPI) traffic.

By using BIG-IP LTM in front of an Outlook Anywhere-enabled Client Access server, you can:

- Offload TLS/SSL processing from the servers to reduce their CPU and memory load
- Simplify certificate management
- Select from a number of advanced load balancing methods to intelligently distribute traffic to the servers.

Critical →

To enable and require SSL for all communications between the Client Access server and the Outlook clients, you must obtain and publish a certificate at the default Web site level. We recommend you purchase your certificate from a third-party certification authority whose certificates are trusted by a wide variety of Web browsers. By default, applications and Web browsers do not trust your root certification authority when you install your own certification authority, such as a BIG-IP self-signed certificate. When a user tries to connect to Microsoft Outlook by using Outlook Anywhere, and the user's computer does not trust the certificate and root Certificate Authority, the connection fails. For more information on this topic, see the following Microsoft TechNet article: http://technet.microsoft.com/en-us/library/aa997703.aspx.

Outlook Anywhere Virtual Server Questions

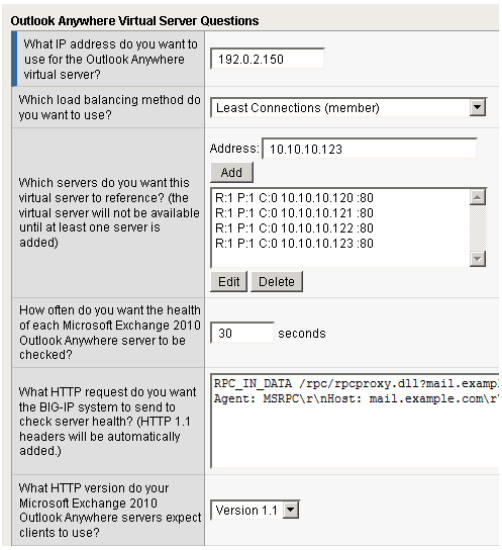

1. **IP address for the virtual server**
   This is the address for Outlook Anywhere. You need an available IP address to use here. *This question does not appear if you chose a single IP address for all HTTP based services.*
2. **Load balancing method**
   While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (member)**.
3. **Address**
   Use the IP address for the Client Access servers running Outlook Anywhere that you entered on the Configuration Worksheet. The template will add the nodes to the appropriate load balancing Pool.
4. **Health Monitor Questions**
   a. *Interval*: Specifies how often the system checks the health of the servers. We recommend the default of 30 seconds.
   b. *HTTP Request*: This is optional but recommended to more accurately monitor the Outlook Anywhere service.
   In the box, type the following string on a single line, replacing the red text with the FQDN of your Outlook Anywhere virtual server.
   RPC_IN_DATA /rpc/rpcproxy.dll?mail.example.com:6001 HTTP/1.1\r\nUser-Agent: MSRPC\r\nHost: mail.example.com\r\n

Critical →

   You must only include a single **\r\n** at the end of the string.
   You must also add a user name and password to this monitor. See *Optional: Modifying the health monitors to add a user name and password on page 27*
   c. *HTTP version*: If you are using the default monitor string, you should leave this at version 1.0. If you are using the custom string above, select **Version 1.1** from the list.

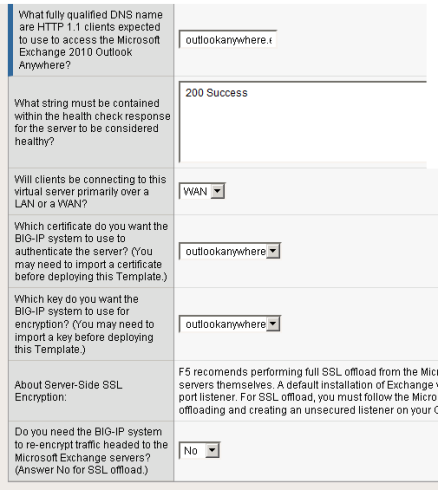

- FQDN: When you select Version 1.1, a new row appears asking for the FQDN for Outlook Anywhere. Type it here.

d. *Monitor response string*: Optional. If you configured the unique HTTP Request we recommended, type the following response string: **`200 Success`**

5. **WAN or LAN**
Specify whether most clients are connecting over a WAN or LAN.

6. **SSL Offload**
Certificate and Key information for Outlook Anywhere. Remember you must have already imported the certificate onto the BIG-IP system before you can select it in the template.

» *Certificate*: Choose the certificate for Outlook Anywhere

» *Key*: Choose the key for Outlook Anywhere.

7. **SSL re-encryption**
While we recommend offloading SSL from the Client Access servers, some organizations require Exchange traffic is encrypted all the way to the Client Access servers, and not offloaded on the BIG-IP. The BIG-IP system has the ability to re-encrypt the traffic before sending it to the Client Access servers.

To get the greatest benefit from SSL offload, select **No** from the list

If you require encrypted traffic from the BIG-IP to the Client Access Servers, select **Yes** from the list, and select the appropriate certificate and key.

Configuring the BIG-IP system for ActiveSync

Exchange ActiveSync is a synchronization protocol based on HTTP and XML that is designed to work over a cellular, wireless Internet or other similar low-bandwidth, high-latency connections. Exchange ActiveSync can synchronize e-mail messages, contacts, calendar, and task data.

By deploying BIG-IP LTM in front of ActiveSync-enabled servers, you gain the following advantages:

- Intelligent load distribution
- SSL/TLS offloading
- Ease of certificate management.

ActiveSync Virtual Server Questions

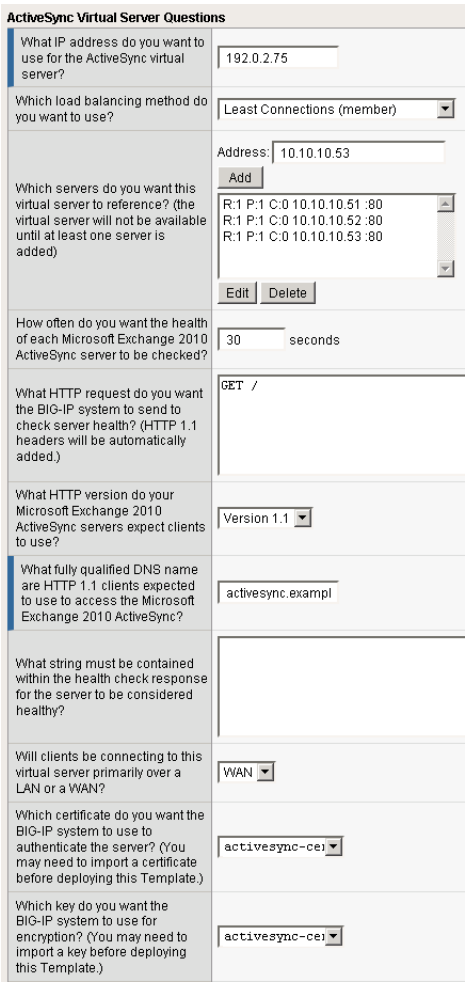

1. **IP address for the virtual server**
   This for the ActiveSync virtual server. You need an available IP address to use here.
   This question does not appear if you chose a single IP address for all HTTP services.
2. **Load balancing method**
   While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (member)**.
3. **Address**
   Use the IP address for the Client Access servers running ActiveSync that you entered on the Preparation Worksheet. The template adds the nodes to the appropriate Pool.
4. **Health Monitor Questions**
   a. *Interval*: How often the system checks server health. We recommend 30 (the default).
   b. *HTTP Request*: This is optional but recommended to more accurately monitor the ActiveSync service. In the box, type the following string on a single line, replacing the red text with the FQDN of your ActiveSync virtual server:
   `OPTIONS /Microsoft-Server-ActiveSync/ HTTP/1.1\r\nHost: mail.example.com\r\n`
   You must also add a user name and password to this monitor. See *Optional: Modifying the health monitors to add a user name and password on page 27*
   c. *HTTP version*: If you are using the default monitor string, you should leave this at version 1.0. If you are using the custom string above, select **Version 1.1** from the list.
      - FQDN: When you select Version 1.1, a new row appears asking for the FQDN for ActiveSync. Type it here.
   d. *Monitor response string*: Optional. If you configured the unique HTTP Request we recommended, type the following response string:
   `MS-ASProtocolCommands: Sync,SendMail,SmartForward,SmartReply,GetAttachment,GetHierarchy,CreateCollection,DeleteCollection,MoveCollection,FolderSync`.
5. **WAN or LAN**
   Specify whether most clients are connecting over a WAN or LAN. We recommend WAN.
6. **SSL Offload**
   Certificate and Key information for ActiveSync. Remember you must have already imported the certificate onto the BIG-IP system before you can select it in the template.
   » *Certificate*: Choose the certificate for ActiveSync
   » *Key*: Choose the key for ActiveSync.

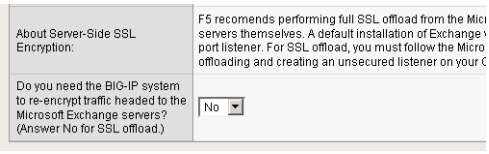

7. **SSL re-encryption**
   While we recommend offloading SSL from the Client Access servers, some organizations require Exchange traffic is encrypted all the way to the Client Access servers, and not offloaded on the BIG-IP. The BIG-IP can re-encrypt the traffic before sending it to the Client Access servers. To get the greatest benefit from SSL offload, select **No**.
   If you require encrypted traffic from the BIG-IP to the Client Access Servers, select **Yes** from the list, and select the appropriate certificate and key.

Be sure to see *Replacing the SSL profile if users are having trouble with iOS and ActiveSync on page 30.*

Configuring the BIG-IP system for Autodiscover

The Autodiscover service provides automatic configuration information to recent versions of Outlook and some mobile clients.

Critical → ***Autodiscover will not work unless you follow the guidelines found at*** *http://technet.microsoft.com/en-us/library/bb124251.aspx*.

We recommend obtaining an SSL certificate that contains three Subject Alternative Names: <domainname>, autodiscover.<domainname>, and the FQDN that you've associated with your Autodiscover service (and that you might be using for other services). For instance, the three Subject Alternative names in a certificate used by the examples in this document would be example.com, autodiscover.example.com, and outlook.example.com. For more information, see *http://technet.microsoft.com/en-us/library/aa995942%28EXCHG.140%29.aspx*.

For more information on Subject Alternative Name certificates, see the *Subject Alternative Name (SAN) SSL Certificates on page 79.*

Autodiscover Virtual Server Questions

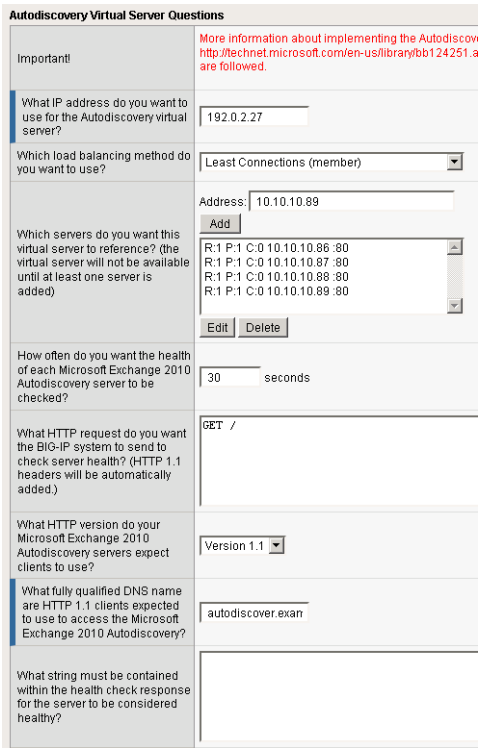

1. **IP address for the virtual server**
   This is the address clients use to access Autodiscover (or a FQDN will resolve to this address). You need an available address to use here.
   This question does not appear if you chose a single IP address for all HTTP based services.
2. **Load balancing method**
   While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (member)**.
3. **Address**
   Use the IP address for the Client Access servers running Autodiscover. The template will add the nodes to the appropriate load balancing Pool.
4. **Health Monitor Questions**
   a. *Interval*: Specifies how often the system checks the health of the servers. We recommend the default of 30 seconds.
   b. *HTTP Request*: This is optional. You can configure the template to retrieve a specific page by typing the path here. In our example, we leave the default.
   c. *HTTP version*: Unless the majority of your users are using HTTP 1.0 (not common), we recommend selecting **Version 1.1** from the list.
      - FQDN: When you select Version 1.1, a new row appears asking for the FQDN the clients use to access Autodiscover. Type it here.
   d. *Monitor response string*: Optional. If you configured a unique HTTP Request, this is where you enter the expected response. In our example, we leave the default.

Tip → *For information on creating a script-based monitor that can more accurately determine if the Autodiscover service is running and returning the correct results, see Optional: Creating an EAV monitor for Autodiscover on page 28.*

5. **WAN or LAN**
   Specify whether most clients are connecting over a WAN or LAN. We recommend selecting WAN (the default).

6. **SSL Offload**
   Certificate and Key information for Autodiscover. Remember you must have already imported the certificate onto the BIG-IP before you can select it in the template.

   » *Certificate*: Choose the certificate for Autodiscover

   » *Key*: Choose the key for Autodiscover.

7. **SSL re-encryption (also known as SSL Bridging)**
   While we recommend offloading SSL from the Client Access servers, some organizations require Exchange traffic is encrypted all the way to the Client Access servers, and not offloaded on the BIG-IP. The BIG-IP system has the ability to re-encrypt the traffic before sending it to the Client Access servers.

   To get the greatest benefit from SSL offload, select **No** from the list

   If you require encrypted traffic from the BIG-IP to the Client Access Servers, select **Yes** from the list.

Configuring the BIG-IP system for RPC Client Access

With Exchange Server 2010, Outlook clients connect using native MAPI to the new RPC Client Access service, which runs on Client Access servers, rather than directly to Mailbox servers. You can use BIG-IP LTM to intelligently balance those incoming MAPI connections.

Unlike most of the other Client Access server roles, the RPC Client Access service does not support offloading of encryption to the BIG-IP LTM.

Note → *Because the RPC Client Access Service requires the BIG-IP LTM to pass traffic to the Client Access servers on a large number of ports, we recommend that you use a firewall to permit only internal networks to access the RPC Client Access virtual server IP address.*

Prerequisites

You should refer to Microsoft documentation regarding the configuration of the RPC Client Access service and mailboxes for each site. However, at a minimum you will need to complete the following steps.

1. In the Microsoft Exchange Management Shell, create a new Client Access Array and associate it with the same FQDN that you will be using. In our example, we type:

   ```
   New-ClientAccessArray -Name "Array01" -FQDN outlook.example.com
   ```

   You should use a name and FQDN appropriate to your deployment.

   Tip → *You can quickly determine your site name using the PowerShell cmdlet* **Get-ADSite**

2. You must modify the attributes of any pre-existing mailbox databases to use the new array. In our example, for existing Mailbox Database 1059170712, we type:

   ```
   Set-MailboxDatabase "Mailbox Database 1059170712" -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.

Important → *After completing the template, be sure to see Modifying the template configuration on page 22 for an important change to the pool.*

RPC Client Access Virtual Server Questions

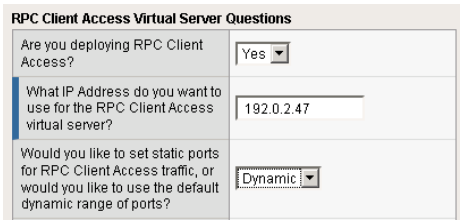

1. **IP address for the virtual server**
   This is the address for the RPC Client Access virtual server. You need an available IP address to use here.
2. **Dynamic port range or static ports**
   You must decide whether you want to specify static ports or a dynamic port range:
   - *Dynamic range*: The system creates one virtual server with dynamic ports for MAPI, Address Book, and the Referral Service. F5 recommends static port assignments for MAPI PRC in Exchange (see the technet links below).
   - *Static Ports*: Additional fields appear. You must specify the ports you want to use for MAPI, Address Book, and the Referral Service. The system creates four virtual servers in this case; one for each port, and one for the RPC Endpoint mapper.
     More information about setting static ports for services can be found at *http://social.technet.microsoft.com/wiki/contents/articles/configuring-static-rpc-ports-on-an-exchange-2010-client-access-server.aspx* and *http://technet.microsoft.com/en-us/library/ee332317.aspx*

If you choose Static Ports, you can type an arbitrary port number for the Referral Service. This virtual server is created by the template, but is unnecessary. See *Required: Modifying the RPC Client Access Configuration on page 22* for instructions on deleting the objects associated with the Referral Service.

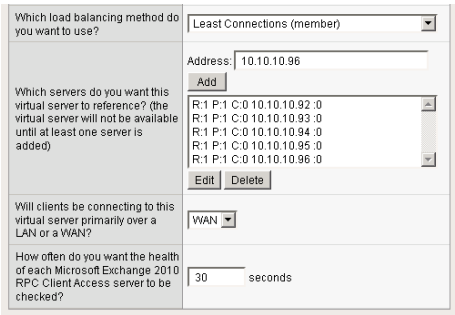

3. **Load balancing method**
   While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (member)**.
4. **Address**
   Use the IP addresses for the Client Access servers running RPC Client Access. The template will add the nodes to the appropriate load balancing Pool.
5. **WAN or LAN**
   Specify whether most clients are connecting over a WAN or LAN.
6. **Health Check interval**
   Specifies how often the system checks the health of the servers. We recommend the default of 30 seconds.

Remember to see *Modifying the template configuration on page 22* for an important change to the pool.

Configuring the BIG-IP system for POP3

POP3 enables a variety of clients to connect to the Exchange server. These include Outlook, Outlook Express, and third-party clients such as Eudora or Mozilla Thunderbird.

The instructions in this guide detail configuring the BIG-IP LTM to serve the secure version of POP3, known as POP3S, perform all SSL/TLS processing, and forward unencrypted traffic to standard ports on the Client Access Servers. Alternate configurations, such as simply forwarding the encrypted traffic, are not included in the template.

For more information about how to manage POP3 in Exchange 2010, see *Understanding POP3 and IMAP4 on Microsoft TechNet* at
http://technet.microsoft.com/en-us/library/bb124107%28EXCHG.140%29.aspx

Prerequisites
By default, the Exchange 2010 POP3 service requires encrypted connections. Because you will be using the BIG-IP LTM to process all secure connections, you must first change the default setting on each Client Access server. You can either change the default setting from the Exchange Management Console or the Management Shell.

To change the default setting using the Exchange Management Console

1. Expand **Server Configuration**, then **Client Access**.
2. In the list of Client Access servers, select a server to which you will be sending POP3 traffic.
3. Select the POP3 protocol, right-click, and then select **Properties**.
4. On the Authentication tab, change the setting to one of the plain text login methods (**Basic** or **Integrated Windows**) as appropriate for your environment and clients.
5. Click **OK**.
6. Restart the POP3 service on the Client Access server.
7. Repeat for each of the Client Access servers to which you will be sending POP3 connections.

To change the default setting using the Exchange Management Shell

1. Type one of the following commands, substituting the name of a Client Access server for "servername":

   For Basic authentication:
   ```
   Set-PopSettings -Server "servername" -LoginType PlainTextLogin
   ```

   For Windows Integrated authentication:
   ```
   Set-PopSettings -Server "servername" -LoginType PlainTextAuthentication
   ```

2. Restart the POP3 service on that Client Access server.
3. Repeat for each of the Client Access servers to which you will be sending POP3 connections.

POP3 Virtual Server Questions

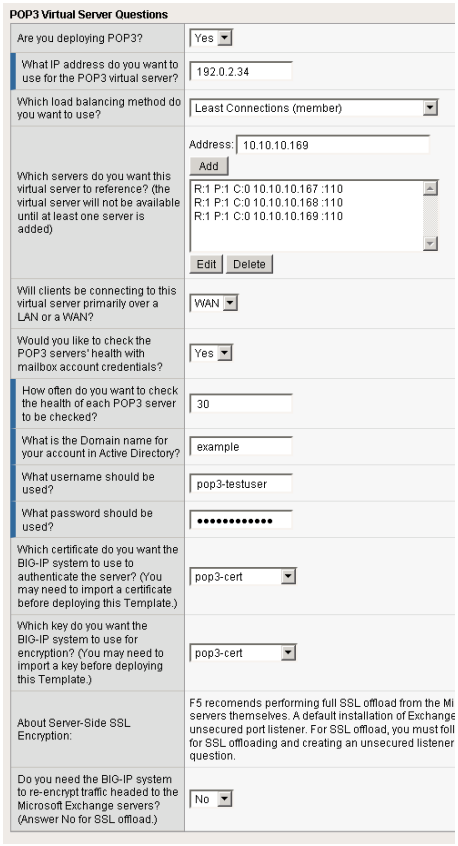

1. **IP address for the virtual server**
   This is the IP address clients use to access the POP3 service.
2. **Load balancing method**
   While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (member)**.
3. **Address**
   Use the IP address for the Client Access servers running POP3. The template will add the nodes to the appropriate load balancing Pool.
4. **WAN or LAN**
   Specify whether most clients are connecting over a WAN or LAN.
5. **Health Monitor Questions**
   a. *Use account credentials to check server health*: The template can be configured created an advanced monitor that attempts to login using a POP3 user account to verify heath status.
      - *No*: By choosing No, the system creates a TCP monitor. Specify the interval of how often the system checks the health of the servers. We recommend the default of 30 seconds.
      - *Yes*: By choosing Yes, the system creates an POP3 monitor. You need to supply the following additional information:
        » *Interval*: Specifies how often the system checks the health of the servers. We recommend the default of 30 seconds.
        » *Domain name*: Type the Domain name for the account in Active Directory.
        » *User name*: Type a valid POP3 user name.
        » *Password*: Type the associated password.
6. **SSL Offload**
   Certificate and Key information for POP3. Remember you must have already imported the certificate onto the BIG-IP system before you can select it in the template.
   » *Certificate*: Choose the certificate for POP3
   » *Key*: Choose the key for POP3.
7. **SSL re-encryption**
   While we recommend offloading SSL from the Client Access servers, some organizations require Exchange traffic is encrypted all the way to the Client Access servers, and not offloaded on the BIG-IP. The BIG-IP system has the ability to re-encrypt the traffic before sending it to the Client Access servers.

   To get the greatest benefit from SSL offload, select **No** from the list

   If you require encrypted traffic from the BIG-IP to the Client Access Servers, select **Yes** from the list.

Configuring the BIG-IP system for IMAP4

As with POP3, IMAP4 enables a variety of clients to connect to the Exchange server.

The instructions detail configuring the BIG-IP LTM to serve the secure version of IMAP4, known as IMAP4S, perform all SSL/TLS processing, and forward unencrypted traffic to standard ports on the Client Access Servers. Alternate configurations, such as simply forwarding the encrypted traffic, are not included in the template.

For more information about how to manage IMAP4 in Exchange 2010, see *Understanding POP3 and IMAP4 on Microsoft TechNet* at
http://technet.microsoft.com/en-us/library/bb124107%28EXCHG.140%29.aspx

Prerequisites

By default, the Exchange 2010 IMAP4 service requires encrypted connections. Because you will be using the BIG-IP LTM to process all secure connections, you must first change the default setting on each Client Access server. You can either change the default setting from the Exchange Management Console or the Management Shell.

To change the default setting using the Exchange Management Console

1. Expand **Server Configuration**, then **Client Access**.
2. In the list of Client Access servers, select a server to which you will be sending IMAP4 traffic.
3. Select the IMAP4 protocol, right-click, and then select **Properties**.
4. On the Authentication tab, change the setting to one of the plain text login methods (**Basic** or **Integrated Windows**) as appropriate for your environment and clients.
5. Click **OK**.
6. Restart the IMAP4 service on the Client Access server.
7. Repeat for each of the Client Access servers to which you will be sending IMAP4 connections.

To change the default setting using the Exchange Management Shell

1. Type one of the following commands, substituting the name of a Client Access server for "servername":

   For Basic authentication:

   ```
   Set-ImapSettings -Server "servername" -LoginType PlainTextLogin
   ```

   For Windows Integrated authentication:

   ```
   Set-ImapSettings -Server "servername" -LoginType PlainTextAuthentication
   ```

2. Restart the IMAP4 service on that Client Access server.
3. Repeat for each of the Client Access servers to which you will be sending IMAP4 connections.

IMAP4 Virtual Server Questions

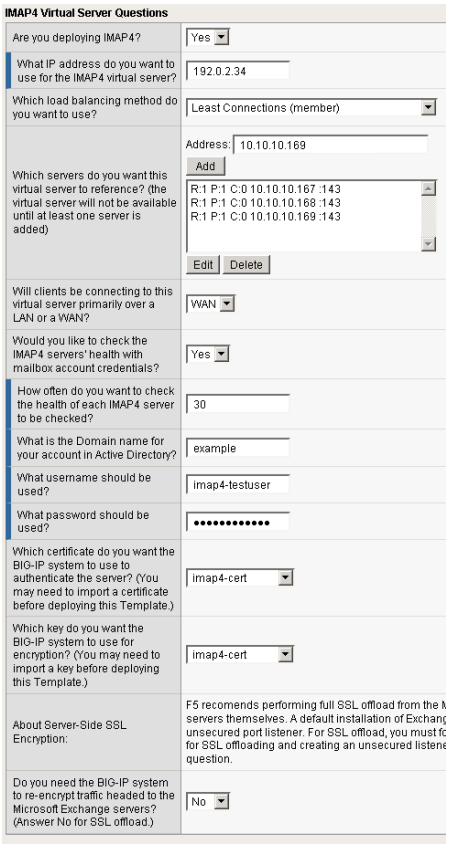

1. **IP address for the virtual server**
   This is the IP address clients use to access the IMAP4 service.
2. **Load balancing method**
   While you can choose any of the load balancing methods from the list, we recommend the default, **Least Connections (member)**.
3. **Address**
   Use the IP address for the Client Access servers running POP3. The template will add the nodes to the appropriate load balancing Pool.
4. **WAN or LAN**
   Specify whether most clients are connecting over a WAN or LAN.
5. **Health Monitor Questions**
   a. *Use account credentials to check server health*: The template can be configured created an advanced monitor that attempts to login using a IMAP4 user account to verify heath status.
      - *No*: By choosing No, the system creates a TCP monitor. Specify the interval of how often the system checks the health of the servers. We recommend the default of 30 seconds.
      - *Yes*: By choosing Yes, the system creates an IMAP monitor. You need to supply the following additional information:
        » *Interval*: Specifies how often the system checks the health of the servers. We recommend the default of 30 seconds.
        » *Domain name*: Type the Domain name for the account in Active Directory.
        » *User name*: Type a valid IMAP4 user name.
        » *Password*: Type the associated password.
6. **SSL Offload**
   Certificate and Key information for IMAP4. Remember you must have already imported the certificate onto the BIG-IP system before you can select it in the template.
   » *Certificate*: Choose the certificate for IMAP4
   » *Key*: Choose the key for IMAP4.
7. **SSL re-encryption**
   While we recommend offloading SSL from the Client Access servers, some organizations require Exchange traffic is encrypted all the way to the Client Access servers, and not offloaded on the BIG-IP. The BIG-IP system has the ability to re-encrypt the traffic before sending it to the Client Access servers.

   To get the greatest benefit from SSL offload, select **No** from the list

   If you require encrypted traffic from the BIG-IP to the Client Access Servers, select **Yes** from the list.

Modifying the 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.

Required: Modifying the RPC Client Access Configuration

If you configured the template for RPC Client Access, there are changes required for the RPC Client Access Pool(s) and profiles that are not yet a part of the template.

In the first procedure, we set the Pool option **Action on Service Down** to **Reject**. This ensures that end users are not sent to an unavailable node in the event of a failure. If you chose Static Ports for RPC Client Access traffic, you must make this change on all RPC Client Access pools.

1. On the Main tab, expand **Local Traffic** and then click **Pools**.
2. From the **Pool** list, select the name of the RPC Client Access pool that was created by the template. By default, this is **my_Exchange_2010__rpc_pool**.
3. In the **Action on Service Down** row, click the **Custom** box, and select **Reject** from the list.
4. Click the **Update** button.
5. If you chose to use Static Ports for RPC Client Access, repeat this procedure for the additional RPC Client Access pools (**my_Exchange_2010__rpc_address_book_pool**, **my_Exchange_2010__rpc_mapi_pool**, and **my_Exchange_2010__rpc_referral_service_pool** by default).

Next, we set a timeout value for the TCP and persistence profiles, and enable Match Across Services and Match Across Virtual Servers on the persistence profile only.

1. Expand **Local Traffic**, click **Profiles**, and then on the Menu bar, click **Persistence**.
2. From the **Persistence** list, select the name of the RPC Client Access profile that was created by the template. By default, this is **my_Exchange_2010__rpc_persist_profile**.
3. In the **Match Across Services** row, check the **Custom** box and then check the box to enable Match Across Services.
4. In the **Match Across Virtual Servers** row, check the **Custom** box and then check the box to enable Match Across Virtual Servers.
5. In the **Idle Timeout** row, check the **Custom** box, and then type **7200** in the box.
6. Click the **Update** button.
7. Under **Local Traffic**, click **Profiles**. On the Menu bar, select **Protocol** and then click **TCP**.
8. From the list, select the name of the RPC Client Access LAN optimized TCP profile that was created by the template. By default, this is **my_Exchange_2010__rpc_lan-optimized_tcp_profile**.
9. In the **Idle Timeout** row, check the **Custom** box, and then type **7200** in the box.
10. Click **Update**.
11. Repeat steps 6-9 for the WAN optimized TCP profile (**my_Exchange_2010__rpc_wan-optimized_tcp_profile** by default).

The next task is to add the RPC 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.

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, by default this is **my_Exchange_2010__single_https_virtual_server**. If you created separate virtual servers and deployed the template for Outlook Anywhere, by default this is **my_Exchange_2010__oa_https_virtual_server.**

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**.

If you configured RPC Client Access to use Static Ports only, we delete all BIG-IP objects associated with the Referral Service, as they are unnecessary.

1. On the Main tab, expand **Local Traffic** and then click **Virtual Servers**.
2. Click a check in the box for the virtual server created by the template for the Referral Service. By default, this is **my_Exchange_2010__rpc_referral_service_virtual_server**.
3. Click the **Delete** button, and click **Delete** again to confirm the deletion.
4. On the Main tab, under **Local Traffic**, click **Pools**.
5. Click a check in the box of the Referral Service pool that was created by the template. By default, this is **my_Exchange_2010__rpc_referral_service_pool**.
6. Click the **Delete** button, and click **Delete** again to confirm the deletion.
7. On the Main tab, under **Local Traffic**, click **Monitors**.
8. Click a check in the box of the Referral Service monitor that was created by the template. By default, this is **my_Exchange_2010__rpc_referral_service_monitor**.
9. Click the **Delete** button, and click **Delete** again to confirm the deletion.

Required: Modifying the HTTP profile

The application template currently incorrectly sets three of the options incorrectly on the HTTP profile for the single virtual server (and on the OWA specific virtual server if you configured separate virtual servers). There are also specific URIs to Exclude from caching if you are deploying OWA.

To modify the HTTP profile

1. On the Main tab, expand **Local Traffic** and then click **Profiles.**
   The HTTP profiles page opens.
2. Locate the HTTP profile used by the single virtual server, or the OWA-specific virtual server. This profile does not use the custom prefix, but instead starts with **microsoft_exchange_2010_** and ends with **_shared_http**. In our example, we click **microsoft_exchange_2010_https_http_wan-optimized-compression_shared_http**.
3. *If you selected **WAN** when asked where clients are primarily connecting from while configuring the template for the single virtual server or any of the HTTP-based services:*
   a. From the **Parent Profile** list, select **http-wan-optimized-compression-caching**.
4. From the **Redirect Rewrite** list, select **All**.
5. From the **Keep Accept Encoding** list, clear the box to ***disable*** Keep Accept Encoding.
6. From the **Insert X-Forwarded-For** list, select **Enabled**.
7. In the **RAM Cache** section, check the **Custom** box if necessary, and then from the **URI Caching** list, select **URI List**.

8. In the **URI** box, type the following URIs, clicking the **Exclude** button after each:
   uglobal.js
   /owa/ev.owa
   oab.xml
9. Click **Update**.

Required: Modifying the persistence iRule

F5 Networks has created new persistence iRules if you are using a single virtual server for all HTTP-based Client Access services, or if you are using separate virtual servers and deploying Outlook Anywhere, that offers functionality not included in the iRules created by the template.

Use the procedure applicable for your configuration.

Important → *If you plan on using Edge Gateway or APM (as described later in this guide),* ***do not*** *modify this iRule. The configuration for Edge Gateway and APM contains specific iRules for each scenario.*

New iRule if you are using a single virtual server for all HTTP-based Client Access services
Use the following procedure if you chose a single virtual server for all HTTP-based CAS services.

To update the persistence iRule

1. On the Main tab, expand **Local Traffic** and then click **iRules**.
2. Click a check in the box for the combined iRule created by the template. By default, this is **my_Exchange_2010__single_Persist_IRule_irule**.
3. In the **Definition** section, copy and paste the iRule found in *Persistence iRule if using a single virtual server for all HTTP-based services on page 75*. It is also available at: http://www.f5.com/solution-center/deployment-guides/files/exchange-persist.zip

   Critical → *Whether you download or copy and paste the iRule ,* ***you must change all Pool names in the iRule to match the pool names in your configuration***.
4. Click the **Update** button.

New iRule if you are using separate virtual servers and deployed Outlook Anywhere
Use following procedure to update the persistence iRule if you chose separate virtual servers and deployed Outlook Anywhere.

To update the persistence iRule

1. On the Main tab, expand **Local Traffic** and then click **iRules**.
2. Click a check in the box for the combined iRule created by the template. By default, this is **my_Exchange_2010__oaPersistIRule**.
3. In the **Definition** section, copy and paste the iRule found in *Outlook Anywhere persistence iRule if using separate pools AND virtual servers on page 78.*
4. Click the **Update** button.

Required: Modify the Cookie persistence profile timeout value

The 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

2. Open a command prompt on the BIG-IP system.

3. If necessary, enter the TMSH shell by typing **tmsh**.
4. At the tmos prompt, use the following command syntax:
   modify /ltm persistence cookie <profile_name> timeout 0

   Where <app_name> is the name you gave the template, and profile name is the name of the Cookie persistence profile created by the template. By default this is **my_Exchange_2010__cookie_persistence_profile**.
   In our example, we type **modify /ltm persistence cookie my_Exchange2010__cookie_persistence_profile timeout 0**

Required: Modify the OneConnect profile

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. Additionally, if you are deploying the template for POP3 and/or IMAP4, you must remove the OneConnect profile from the virtual server(s).

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 template. By default, if you chose a single virtual server, this is **my_Exchange_2010__single_one_connect_profile.** If you chose separate virtual servers, this is **my_Exchange_2010__<service abbreviation>_one_connect_profile**, where <service abbreviation> is **owa**, **ad**, and/or **as**.
4. In the **Source Mask** box type **255.255.255.255** in the box. When you are finished, click the **Update** button.
5. Repeat this procedure to modify the OneConnect profile for the other Client Access HTTP-based services (OWA, Autodiscover, and ActiveSync).

To remove the OneConnect profile from the POP3 and IMAP4 virtual servers

If you deployed the template for POP3 and/or IMAP4 only:

1. On the Main tab, expand **Local Traffic** and then click **Virtual Servers.**
2. Click the name of the POP3 or IMAP4 virtual server created by the template. By default for POP3, this is **my_Exchange_2010__pop3_virtual_server**, and for IMAP4, this is **my_Exchange_2010__imap4_virtual_server**.
3. From the **OneConnect Profile** list, select **None**.
4. Click the **Update** button.
5. If you deployed both POP3 and IMAP4, repeat this procedure for the other service.

Required: Modifying the Append iRule for OWA

The Append iRule for Outlook Web App created by the template contains a trailing slash (/owa/) that may result in the browser not appending /owa/ after the initial request in a single session (for example, if a user browses to Outlook Web App, then to another site, and then back to Outlook Web App). In this case, you must modify the iRule so that the BIG-IP LTM performs a redirect rather than appending anything.

To modify the Append iRule

2. On the Main tab, expand **Local Traffic** and then click **iRules.**
3. Click the name of the OWA Append iRule created by the template. This iRule is preceded by

the name you gave the template, followed by **_owa_append_irule**.

4. Replace the existing iRule code with the following iRule

```
when HTTP_REQUEST {
    if { ([HTTP::uri] == "/") } {
        HTTP::redirect https://[HTTP::host]/owa/
    }
}
```

5. Click the **Update** button.

Using a SNAT Pool if you expect more than 6,000 users per Client Access server

The BIG-IP system can create roughly 64,000 connections per SNAT address. 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, if you expect more than approximately 6,000 users per Client Access server, you should use a SNAT Pool, and not SNAT Automap as configured by the template. The SNAT pool should contain an IP address for each 6,000 users you expect, or fraction thereof. For example, if you anticipate 15,000 users for each Client Access Server, you would configure a SNAT Pool with at least three IP addresses.

To configure a SNAT Pool, from the BIG-IP Configuration utility, expand **Local Traffic**, and then click **SNATs**. On the Menu bar, click **SNAT Pool List**. Click the **Create** button and configure the SNAT Pool as applicable for your configuration. For more information on configuring SNAT Pools, see the Online Help or the BIG-IP LTM documentation.

Creating an iRule when using a SNAT pool

If using a SNAT Pool, multiple connections from a single client are split between multiple source IP addresses by default. As a result, some services, such as the Outlook Client and Blackberry® Enterprise Server that use multiple connections to the RPC Client Access service, may not function properly without the following iRule.

To create the iRule, from the BIG-IP Configuration utility, expand **Local Traffic**, and then click **iRules**. Click the **Create** button, give the iRule a name, and then use the following code in the **Definition** section. Do not include the line numbers on the left, and modify the IP addresses in the following example to SNAT Pool IP addresses you defined when configuring the SNAT Pool.

```
when RULE_INIT {
  # Use a local array to configure SNAT addresses.
  # These addresses must be defined in a SNAT pool
  # In this example, we use three addresses. Replace
  # these with the IP addresses used in your SNAT Pool.
  # Follow the pattern of the existing addresses to add more than three.

  set static::snat_ips(0) 10.0.0.1
  set static::snat_ips(1) 10.0.0.2
  set static::snat_ips(2) 10.0.0.3
}

when CLIENT_ACCEPTED {
  # Calculate the crc32 checksum of the client IP.
  # Use the modulo of the checksum and the number of SNAT IPs in the array
  # to select a SNAT IP address.

  snat $static::snat_ips([expr {[crc32 [IP::client_addr]] % [array size static::snat_ips]}])

}
```

After creating the SNAT pool and the iRule, you must modify the virtual server(s) created by the template to use the SNAT Pool and the iRule.

From the Configuration utility, under Local Traffic, click **Virtual Servers**. Click the first virtual server created by the template. From the **SNAT Pool** list, select the SNAT Pool you created, and then click the **Update** button. If you created the iRule, on the Menu bar, click **Resources**. In the iRules section, click **Manage**, add the iRule, and then click **Finished**. Repeat for each virtual server created by the template (with the exception of the redirect virtual server on port 80).

Optional: Modifying the health monitors to add a user name and password

If you configured any of the health monitors with the optional Send and Receive Strings, you must modify the monitors to include a user name and password of a valid Exchange account. This includes the monitors for OWA if using Basic or Basic and Windows Integrated Authentication and not Forms-based authentication (default), Outlook Anywhere, and ActiveSync.

Important

See *Important note about BIG-IP health monitors that use Exchange server accounts on page 30.*

To modify the monitors

1. On the Main tab, expand **Local Traffic** and then click **Monitors**.
2. From the Monitor list, select the name of the applicable monitor that was created by the template. For example, for Outlook Web App this is **my_Exchange_2010__owa_monitor**.
3. In the **Send String** box, make sure the appropriate string is present.
4. In the **Receive String** box, make sure the appropriate string is present.
5. In the **User Name** box, type the user name of the Exchange account for this monitor.
6. In the **Password** box, type the associated password.
7. Click the **Update** button.

Required: Adding the ActiveSync persist iRule if using separate virtual servers

If you are deploying ActiveSync and using separate virtual servers on a BIG-IP system behind a NAT or other address aggregating device, use this iRule to ensure even distribution of client connections. This iRule is only necessary if you deployed separate virtual servers and ActiveSync.

To add the iRule

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

```
1  when HTTP_REQUEST {
2      if { [HTTP::header exists "Authorization"] } {
3          set as_key [sha256 [HTTP::header "Authorization"]]
4          persist uie $as_key 7200
5      } else {
6          persist source_addr
7      }
8  }
```

5. Click **Finished**. Now you attach the iRule to the ActiveSync virtual server.
6. On the Main tab, expand **Local Traffic** and then click **Virtual Servers.**
7. Click the name of the ActiveSync virtual server created by the template. By default this is **my_Exchange_2010__as_https_virtual_server**.

8. On the Menu bar, click **Resources**.
9. In the iRules section, click **Manage**.
10. From the **Available** list, select the name of the iRule you just created, and then click the Add (**<<**) button to enable it.
11. Click the **Update** button.

Modifying the IIS authentication token timeout value

The F5 recommends you configure most Exchange monitors to check service health every 30 seconds. However, to reduce traffic between the Exchange server and domain controllers, IIS virtual directories configured to use Basic authentication cache authentication tokens for up to 15 minutes before re-authenticating the user with Active Directory. This may result in the BIG-IP pool members for these services being marked UP incorrectly while Basic authentication tokens are cached.

You can decrease the length of or disable this token caching period by editing the registry on the Exchange server. The length of time configured for the token cache combined with the timeout value of the monitor will determine how long it will take until a resource is marked down. For example, setting a token cache period of 60 seconds, combined with a monitor using a timeout value of 91 seconds, will result in a resource being marked down after 151 seconds.

For instructions on modifying the registry, see the following Microsoft article (while this article says IIS 6.0, we tested it on IIS 7.5 with no modifications):

Critical →

Use extreme caution any time you are editing the registry. Contact Microsoft for specific instructions and/or help editing the registry values.
http://www.microsoft.com/technet/prodtechnol/WindowsServer2003/Library/IIS/6b2e7fcd-5fad-4ac8-ac0a-dcfbe771e9e1.mspx

Optional: Creating an EAV monitor for Autodiscover

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 requires a valid Exchange server account and associated mailbox specifically for monitoring purposes.

EAVs cannot be created or modified by the application template. And 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 appropriate script file on each BIG-IP system, create the external monitor that calls the script, then update the load balancing pool to use the monitor.

Important →

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

To download and install the monitor

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

   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.

2. Extract the file and copy the resulting script (autodiscover-monitor.sh or autodiscover-monitor-ssl-bridging.sh) to the appropriate directory on the BIG-IP system.
   If using v10.0.x, the directory is **/usr/bin/monitors/**.
   If using 10.1 or later in the 10.x branch, the directory is **/config/monitors**
3. Change the permissions of the file using the following command (modify path if necessary):
   For SSL Offload: `chmod 755 /config/monitors/autodiscover-monitor.sh`
   For SSL Bridging: `chmod 755 /config/monitors/autodiscover-monitor-ssl-bridging.sh`

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	
Name	User choice.	
Type	**External** (the *Import Settings* field automatically selects External as well)	
Interval	User choice, but we recommend **60**.	
Timeout	User choice, but we recommend **181**.	
External Program	`SSL offload: /usr/bin/monitors/autodiscover-monitor.sh` `SSL bridging: /usr/bin/monitors/autodiscover-monitor-ssl-bridging.sh`	
Variables	*Name*	*Value*
	USER	The account name associated with a mailbox.
	PASSWORD	The password for the account
	DOMAIN	The Windows domain for the account
	EMAIL	The email address for the user mailbox (such as j.smith@example.com)

Important

Be sure to see the *Important note about BIG-IP health monitors that use Exchange server accounts on page 30.*

The final task is to remove the default monitor from the Autodiscover pool and add the new external monitor you created.

To modify the Autodiscover pool

1. On the Main tab, expand **Local Traffic** and then click **Pools**.
2. From the **Pool** list, select the name of the Autodiscover pool that was created by the template. By default, this is **my_Exchange_2010_single_ad_pool**.
3. From the Health Monitor section, in the **Active** box, click the name of the monitor created by the template, and then click the Remove (**>>**) button.
4. From the **Available** box, click the name of the external monitor you just created and click the **Add** button to add it to the **Active** box.
5. Click the **Update** button.

Critical →

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-IP systems, make sure the credentials are not used elsewhere in your organization for anything other than monitoring.

Note that BIG-IP v10.x health monitors do not support NTLM authentication. If you require NTLM authentication support for your monitors, we suggest you upgrade to v11 and use the iApp template to configure the BIG-IP system.

Replacing the SSL profile if users are having trouble with iOS and ActiveSync

If you deployed the iApp template for ActiveSync (or manually configured the BIG-IP system) and iOS devices started showing invalid certificate messages even though the certificates were issued by an appropriate authority, you must modify the Client SSL profile to use a Chain certificate.

Intermediate certificates, also called intermediate certificate chains or chain certificates, are used to help systems which depend on SSL certificates for peer identification. See *http://support.f5.com/kb/en-us/solutions/public/13000/300/sol13302.html* for help on creating an intermediate certificate chain.

To modify the SSL profile

1. On the Main tab, expand **Local Traffic** and then click **Profiles.**
2. On the Menu bar, from the **SSL** menu, select **Client**.
3. Click the name of the Client SSL profile created by the template.
4. Modify the profile to use a chain certificate. See *http://support.f5.com/kb/en-us/solutions/public/13000/300/sol13302.html* for details.
5. Be sure **Secure Renegotiation** is set to **Require**.
6. Click the **Update** button.

Chapter 2

Secure Access to Exchange 2010 Client Access Servers

F5 has three solutions that can be used for secure access to Client Access servers:

- The BIG-IP Edge Gateway - an access solution that brings together SSL VPN remote access, security, application acceleration, and availability services for remote users.
  If you are using the BIG-IP Edge Gateway, you must be on version BIG-IP **10.2.2** or later.
- The BIG-IP Access Policy Manager (APM) - a flexible, high-performance access and security solution. If you are using the BIG-IP APM, you must be on BIG-IP version **10.2.2** or later.
- The FirePass SSL VPN appliance - provides secure remote access to enterprise applications and data for users over any device or network.

For more information on the FirePass Controller, see *www.f5.com/products/firepass*

For more information on the BIG-IP Edge Gateway, see *www.f5.com/products/big-ip/edge-gateway.html*

For more information on the BIG-IP Access Policy Manager, see *www.f5.com/products/big-ip/access-policy-manager.html*

Configuring the BIG-IP Edge Gateway or Access Policy Manager for Client Access servers

Using the F5 Edge Gateway or Access Policy Manager (APM), BIG-IP versions 10.2.2 and later support proxying and authentication of traffic for all four of the HTTPS-based Exchange Client Access services: Outlook Web App (which includes the Exchange Control Panel and Exchange Web Services), Outlook Anywhere, ActiveSync and Autodiscover.

You can use Edge Gateway or APM to support authentication and session management for a variety of clients and usage scenarios. For instance, when deployed in front of Exchange 2010, F5 can authenticate and manage Outlook Web App sessions, Outlook 2007 and Outlook 2010 client sessions, and sessions from clients that use Exchange Web Services (EWS) and the ActiveSync protocol, such as Microsoft Office 2008 and 2011 for Mac, or the Mac OS X-native Mail, Calendar, and Address Book applications.

This chapter covers the following scenarios:

1. An Edge Gateway deployment on a separate BIG-IP than that providing your Exchange 2010 traffic management. 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.
2. An Edge Gateway deployment on a separate BIG-IP than that providing your Exchange 2010 traffic management. 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, and then re-encrypt the traffic before placing it back on the network.
3. A single BIG-IP configured with both APM and LTM modules. The BIG-IP will terminate SSL connections and forward traffic to your Exchange Client Access servers on HTTP port 80.

4. A single BIG-IP configured with both APM and LTM modules. 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.
5. The FirePass SSL VPN for secure remote access to the Client Access servers

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).

Configuration example

In this scenario (the first scenario described above), you have separate internal and external DNS servers; you want your external Outlook users to connect via Outlook Anywhere, and internal users to connect via MAPI.

External users receive an authentication prompt on initial connection. Internal users, if they are using a domain-member Windows client and logged on using their domain credentials, do not.

We assume you have enabled all HTTP services (OWA, OA, Autodiscover, ActiveSync) using a single virtual server, and configured them for SSL Offload as described in this guide.

The following diagram shows the traffic flow in this scenario. This example shows the BIG-IP LTM offloading SSL.

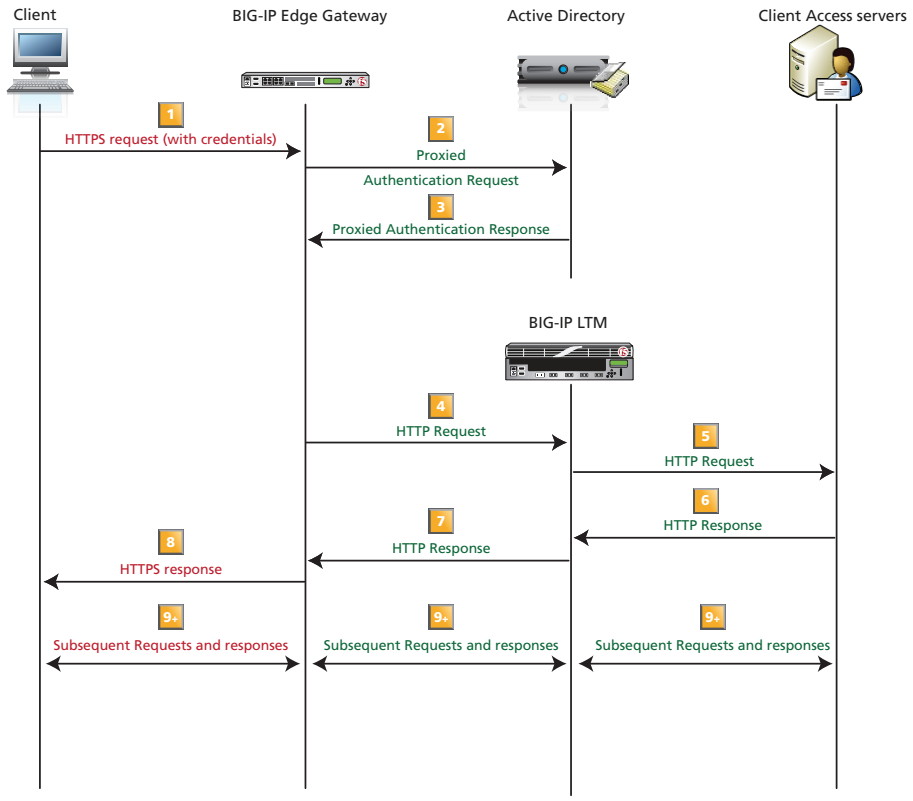

Figure 2.1: *Edge Gateway flow diagram*

You should follow the steps detailed in Chapter 1 in this guide for configuration of BIG-IP LTM using the application template (highly recommended), or build the required objects manually using the configuration tables in *Appendix A: Manual Configuration Tables on page 50*. Once complete, you need to modify some settings both in BIG-IP and on your Exchange Server Client Access servers before performing the Edge Gateway or APM configuration steps. A summary of the required changes or important configuration requirements is listed below.

Prerequisites and configuration notes

- If using the BIG-IP APM or Edge Gateway, you ***must*** be on version 10.2.2 or later.
- If the BIG-IP Edge Gateway is re-encrypting traffic before sending it to the BIG-IP LTM (not recommended), and internal users do not go through the Edge Gateway for RPC Client Access, you ***must*** use a valid certificate (usually SAN-enabled) and not the default, self-signed certificate for the Client SSL profile on the BIG-IP LTM.
  If not re-encrypting traffic, you do not need a third party certificate.

Important →

- For the scenarios in this chapter, you must have used the application template to configure a single virtual server for all HTTP-based services as described earlier in this guide.
- If using a single virtual server for all HTTP 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. The BIG-IP does not display SAN values, but uses these certificates correctly.

Note →

For more information on SAN certificates, see *Subject Alternative Name (SAN) SSL Certificates on page 79*.

- You ***must*** use the Forms-based monitor for Outlook Web App as described in the step 4, Health monitor section, of the Outlook Web App configuration on *page 8*.
- Ensure that your BIG-IP Edge Gateway or BIG-IP APM is configured to use a DNS server that is part of your Active Directory, or that forwards queries to Active Directory.
- *Exchange Server (Client Access Server) settings*

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

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

[2] *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.*

[3] *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 clients to the appropriate virtual server):

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

Your network topology may differ considerably from the example shown.

Note → *You may choose to use separate names for all four HTTP services and the RPC Client Access service (Client Access Array).*

DNS Settings

Record	External DNS	Internal DNS
A Records	owa.example.com: 192.0.2.10 outlook.example.com: 192.0.2.11 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: example.com: 192.0.2.10 autodiscover.example.com: 192.0.2.10	owa.example.com: 192.0.2.10 mapi.example.com: 10.0.0.10 If the SRV record listed below is not used **and** 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: example.com: 192.0.2.10 autodiscover.example.com: 192.0.2.10 To prevent internal users from receiving a password prompt, internal DNS must not have an A record for the FQDN for Outlook Anywhere.
SRV Records	_autodiscover._tcp.example.com: port 443, host 'owa.example.com'	_autodiscover._tcp.example.com: port 443, host 'owa.example.com' (optional; Outlook can use SCP instead. See note above and Further Reading below)

Further reading:

- Summary of SRV records on Wikipedia: *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*

Scenario 1: Configuring the BIG-IP Edge Gateway (SSL offload)

In scenario 1, your Edge Gateway deployment is on a separate BIG-IP than that providing your Exchange 2010 traffic management. SSL (HTTPS, port 443) connections are terminated at the Edge Gateway and forwarded to the BIG-IP LTM and then to your Exchange Client Access servers on HTTP port 80. ***All load balancing and persistence operations occur on the BIG-IP LTM.*** Modify the LTM configuration as described in the following section, and then use the table to configure the Edge Gateway.

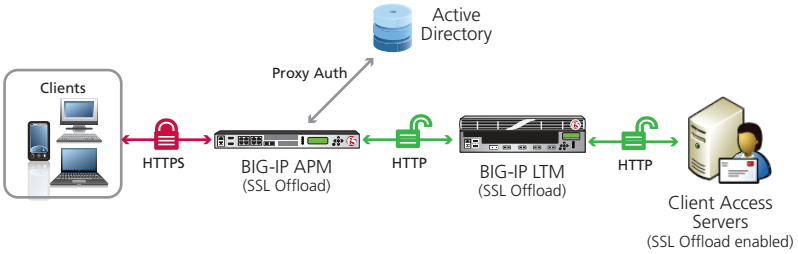

Figure 2.2: *Edge Gateway with SSL Offload configuration example*

For scenario 1, you must complete the following procedures:

- *Creating the iRules on the Edge Gateway and BIG-IP LTM, on this page*
- *Deleting the redirect virtual server and modifying the 443 virtual server on the LTM on page 37*
- *BIG-IP Edge Gateway and APM Configuration on page 45*
- *Configuring the Edge Gateway for Scenarios 1 and 2 on page 47*

Creating the iRules on the Edge Gateway and BIG-IP LTM

The next task is to create the iRules on the BIG-IP LTM for Edge Gateway. The first iRule is necessary for all deployments with Edge Gateway. The second is only necessary if ***all*** traffic goes through the Edge Gateway on the way to the BIG-IP LTM.

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.

```
when ACCESS_ACL_ALLOWED {
    set sessionid [ACCESS::session data get "session.user.sessionid"]
    HTTP::header insert APM_session $sessionid
}
when HTTP_RESPONSE {
    if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate"} {
        ONECONNECT::reuse disable
        ONECONNECT::detach disable
        ## disables NTLM conn pool for connections where OneConnect has been disabled
        NTLM::disable
    }
    ## this command rechunks encoded responses
    if {[HTTP::header exists "Transfer-Encoding"]} {
        HTTP::payload rechunk
    }
}
```

4. Click the **Finished** button.

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

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.*

1. On the Main tab, expand **Local Traffic**, click **iRules**, and then click the **Create** button.
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.

```
when HTTP_REQUEST {

## CHANGE ALL POOL NAMES TO MATCH THOSE IN YOUR ENVIRONMENT.
## OAB and Autodiscover do not require persistence.

      switch -glob -- [string tolower [HTTP::path]] {
          "/microsoft-server-activesync" {
                    pool my_Exchange_2010__single_as_pool
                    COMPRESS::disable
                    ## If you selected LAN when asked from where clients are primarily
                    ## connecting, you MUST remove or comment out the CACHE::disable line
                    CACHE::disable
                    persist uie [HTTP::header "APM_session"] 7200
                    return
                    }
          "/ews*" {
                    pool my_Exchange_2010__single_owa_pool
                    COMPRESS::disable
                    ## If you selected LAN when asked from where clients are primarily
                    ## connecting, you MUST remove or comment out the CACHE::disable line
                    CACHE::disable
                    persist uie [HTTP::header "APM_session"] 7200
                    return
                    }
          "/ecp*" {
                    pool my_Exchange_2010__single_owa_pool
                    persist uie [HTTP::header "APM_session"] 7200
                    return
                    }
          "/oab*" {
                    pool my_Exchange_2010__single_owa_pool
                    return
                    }
          "/rpc/rpcproxy.dll" {
                    pool my_Exchange_2010__single_oa_pool
                    COMPRESS::disable
                    ## If you selected LAN when asked from where clients are primarily
                    ## connecting, you MUST remove or comment out the CACHE::disable line
                    CACHE::disable
                    persist uie [HTTP::header "APM_session"] 7200
                    return
                    }

          "/autodiscover*" {
                    pool my_Exchange_2010__single_ad_pool
                    return
                    }
          default {
                    ## This final section takes all traffic that has not otherwise
                    ## been accounted for and sends it to the pool for Outlook Web
                    ## App
                    pool my_Exchange_2010__single_owa_pool
                    persist uie [HTTP::header "APM_session"] 7200
                    }
      }
}
when HTTP_RESPONSE {
    if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate"} {
        ONECONNECT::reuse disable
        ONECONNECT::detach disable
        ## disables NTLM conn pool for connections where OneConnect has been disabled
        NTLM::disable
    }
## this command rechunks encoded responses
    if {[HTTP::header exists "Transfer-Encoding"]} {
        HTTP::payload rechunk
    }
}
```

Critical

You must change the pool names in this iRule to match the names of the pools in your configuration.

This iRule continues on the following page.

4. Click the **Finished** button.

Deleting the redirect virtual server and modifying the 443 virtual server on the LTM
For this scenario, you must first delete the port 80 redirect virtual server on the BIG-IP LTM created by the template. After deleting the virtual server, you modify the remaining virtual server for the Client Access services from port 443 to port 80, remove the Client SSL profile and add the iRule if all traffic is going through the BIG-IP Edge Gateway.

To delete the port 80 virtual server

1. On the Main tab of the BIG-IP LTM, expand **Local Traffic** and then click **Virtual Servers**.
2. From the **Virtual Server** list, check the box for the single virtual server for the Client Access HTTP-based services on port 80 created by the template. By default, this is **my_Exchange_2010__single_virtual_server**. Make sure you do ***not*** delete the HTTPS virtual (*my_Exchange_2010__single_https_virtual_server*).
3. Click the **Delete** button, and then confirm the deletion.

Next, we modify the HTTPS virtual server port from port 443 to port 80 and remove the SSL profile.

To modify the port 443 virtual server on the BIG-IP LTM

1. On the Main tab of the BIG-IP LTM, expand **Local Traffic** and then click **Virtual Servers**.
2. From the **Virtual Server** list, click the name of the single virtual server for the Client Access HTTP-based services on port 443 created by the template. By default, this is **my_Exchange_2010__single_https_virtual_server**.
3. In the **Service Port** box, type **80**, or select **HTTP** from the list.
4. From the **SSL Profile (Client)** list, select **None**.
5. Click **Update**.
6. If you did not create the iRule for all traffic going through the Edge Gateway, no further modifications are necessary, continue with the configuration table on page 45.
7. If you created the iRule for all traffic going through the Edge Gateway only: On the Menu bar, click **Resources**.
8. In the iRules section, click **Manage**.
9. From the **Enabled** box, select the Persistence iRule create by the template and then click the Remove (**>>**) button. By default this is **my_Exchange_2010__single_Persist_irule**.
10. From the **Available** list, click the name of the persistence iRule you just created and then click the Add (**<<**) button to move it to the Enabled box.
11. Click the **Finished** button.

Creating the iRule to terminate inactive sessions
APM and Edge Gateway sessions can remain active after users have either manually logged out of OWA or the OWA session has timed out due to user inactivity. This iRule checks the OWA session status and terminates the associated APM session if applicable.

To create the inactive sessions iRule, use the procedure *Creating the iRule to terminate inactive sessions on page 42*.

After completing these changes, continue with the configuration table on page 45.

Scenario 2: Configuring the BIG-IP Edge Gateway (SSL Bridging)

In this scenario, the Edge Gateway deployment is on a separate BIG-IP than that providing your Exchange 2010 traffic management. Both the BIG-IP Edge Gateway and the BIG-IP LTM perform SSL Bridging; they decrypt SSL traffic in order to process it, and then re-encrypt the traffic before placing it back on the network. ***All load balancing and persistence operations occur on the BIG-IP LTM.***

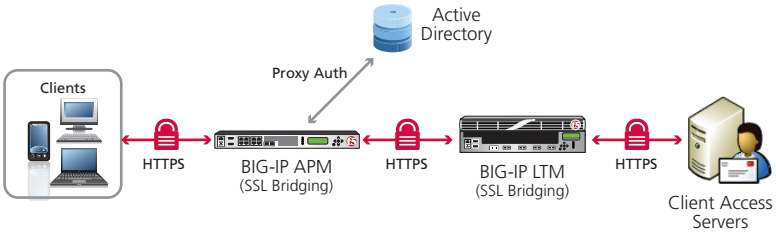

Figure 2.3: *Edge Gateway with SSL Bridging configuration example*

For scenario 2, you must complete the following procedures:

- *Deleting the redirect virtual server, on this page*
- *Creating the new HTTPS monitors, on this page*
- *Creating the iRules for Edge Gateway on page 39*
- *Modifying the virtual server to use the iRule on page 39*
- *BIG-IP Edge Gateway and APM Configuration on page 45*
- *Configuring the Edge Gateway for Scenarios 1 and 2 on page 47*

Deleting the redirect virtual server on the BIG-IP LTM

For this scenario, we recommend you delete the redirect virtual server on the BIG-IP LTM because it is not used.

To delete the port 80 virtual server

1. On the Main tab, expand **Local Traffic** and then click **Virtual Servers**.
2. From the **Virtual Server** list, check the box for the single virtual server for the Client Access HTTP-based services on port 80 created by the template. By default, this is **my_Exchange_2010__single_virtual_server**. Make sure you do ***not*** delete the HTTPS virtual (*my_Exchange_2010__single_https_virtual_server*).
3. Click the **Delete** button, and then confirm the deletion.

Creating new HTTPS monitors on the BIG-IP LTM

Because the template creates HTTP monitors for the Client Access HTTP-based services, for this scenario, you must create new HTTPS monitors for those services.

To create the new monitors, from the Configuration utility, expand **Local Traffic** and then click **Monitors**. Click the **Create** button. Give the monitor a unique name use the following table to create the monitor with the recommended guidance. Repeat for each Client Access service.

CAS Service	Monitor type
Outlook Web App	**HTTPS** parent (see recommended health monitor configuration in Step 4 on *page 8*)
Outlook Anywhere	**HTTPS** parent (see recommended health monitor configuration in Step 4 on *page 10*)
ActiveSync	**HTTPS** parent (see recommended health monitor configuration in Step 4 on *page 12*)
Autodiscover	**HTTPS** parent (see recommended health monitor configuration in Step 4 on *page 14*)

After creating the monitors, you must associate the new monitors with the relevant pool.

To associate the HTTPS monitor with the relevant pool

1. On the Main tab, expand **Local Traffic** and then click **Pools**.
2. From the **Pools** list, click the name of the Outlook Web App pool created by the template. By default, this is **my_Exchange_2010__single_owa_pool**.
3. In the Health Monitors section, from the **Active** box, select the existing health monitor (**my_Exchange_2010__single_owa_monitor** by default) and then click the Remove (**>>**) button.
4. From the **Available** box, select the name of the new HTTPS Outlook Web App monitor you created and then click the Add (**<<**) button to move it to the Active box.
5. Click the **Update** button.
6. Repeat this procedure for each of the HTTP-based Client Access services (_single_ad_pool, _single_as_pool, _single_oa_pool).

Creating the iRules for Edge Gateway

The next task is to create the iRules for Edge Gateway. To create the iRules, use the procedure *Creating the iRules on the Edge Gateway and BIG-IP LTM on page 35*. After creating the iRule, return to this section to modify the virtual server.

Modifying the virtual server to use the iRule

If you created the iRule for all traffic going through the Edge Gateway, you must modify the virtual server to use the iRule you just created.

To modify the virtual server

1. On the Main tab, expand **Local Traffic** and then click **Virtual Servers**.
2. From the **Virtual Server** list, click the name of the single virtual server for the Client Access HTTP-based services on port 443 created by the template. By default, this is **my_Exchange_2010__single_https_virtual_server**.
3. On the Menu bar, click **Resources**.
4. In the iRules section, click **Manage**.
5. From the **Enabled** box, select the Persistence iRule create by the template and then click the Remove (**>>**) button. By default this is **my_Exchange_2010__single_Persist_irule**.
6. From the **Available** list, click the name of the persistence iRule you just created and then click the Add (**<<**) button to move it to the Enabled box.
7. Click the **Finished** button.

Creating the iRule to terminate inactive sessions

APM and Edge Gateway sessions can remain active after users have either manually logged out of OWA or the OWA session has timed out due to user inactivity. This iRule checks the OWA session status and terminates the associated APM session if applicable.

To create the inactive sessions iRule, use the procedure *Creating the iRule to terminate inactive sessions on page 42*.

After completing these changes, continue with the configuration table on page 45.

Scenario 3: Configuring the BIG-IP APM (SSL Offload)

In this scenario, a single BIG-IP configured with both APM and LTM modules. The BIG-IP will terminate SSL connections and forward traffic to your Exchange Client Access servers on HTTP port 80.

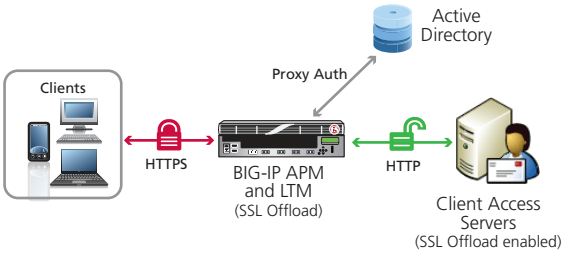

Figure 2.3: *BIG-IP APM with SSL Offload configuration example*

For scenario 3, you must complete the following procedures:

- *Creating the persistence iRule when using BIG-IP APM, on this page*
- *Creating the iRule to terminate inactive sessions on page 42*
- *Modifying the virtual server to use the APM persistence iRule on page 42*
- *BIG-IP Edge Gateway and APM Configuration on page 45*
- *Configuring the BIG-IP APM for Scenarios 3 and 4 on page 48*

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.

Critical

You must change the pool names in this iRule to match the names of the pools in your configuration.

```
## CHANGE ALL POOL NAMES TO MATCH THOSE IN YOUR ENVIRONMENT.
## OAB and Autodiscover do not require persistence.

when ACCESS_ACL_ALLOWED {
     set sessionid [ACCESS::session data get "session.user.sessionid"]

     switch -glob -- [string tolower [HTTP::path]] {
          "/microsoft-server-activesync" {
               pool my_Exchange_2010__single_as_pool
               COMPRESS::disable
               ## If you selected LAN when asked from where clients are primarily
               ## connecting, you MUST remove or comment out the CACHE::disable line
               CACHE::disable
               persist uie $sessionid 7200
               return
          }
          "/ews*" {
               pool my_Exchange_2010__single_owa_pool
               COMPRESS::disable
               ## If you selected LAN when asked from where clients are primarily
               ## connecting, you MUST remove or comment out the CACHE::disable line
               CACHE::disable
               persist uie $sessionid 7200
               return
          }
          "/ecp*" {
               pool my_Exchange_2010__single_owa_pool
               persist uie $sessionid 7200
               return
          }
          "/oab*" {
               pool my_Exchange_2010__single_owa_pool
               return
          }
          "/rpc/rpcproxy.dll" {
               pool my_Exchange_2010__single_oa_pool
               COMPRESS::disable
               ## If you selected LAN when asked from where clients are primarily
               ## connecting, you MUST remove or comment out the CACHE::disable line
               CACHE::disable
               persist uie $sessionid 7200
               return
          }
          "/autodiscover*" {
               pool my_Exchange_2010__single_ad_pool
               return
          }

          default {
          ## This final section takes all traffic that has not otherwise
          ## been accounted for and sends it to the pool for Outlook Web
          ## App
               pool my_Exchange_2010__single_owa_pool
               persist uie $sessionid 7200
          }
     }
}
when HTTP_RESPONSE {
    if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate"} {
        ONECONNECT::reuse disable
        ONECONNECT::detach disable
        ## disables NTLM conn pool for connections where OneConnect has been disabled
        NTLM::disable
    }
    ## this command rechunks encoded responses
    if {[HTTP::header exists "Transfer-Encoding"]} {
        HTTP::payload rechunk
    }
}
```

4. Click **Finished**.

Creating the iRule to terminate inactive sessions

APM and Edge Gateway sessions can remain active after users have either manually logged out of OWA or the OWA session has timed out due to user inactivity. The following iRule checks the OWA session status and terminates the associated APM session if applicable.

To add the APM Edge Gateway session check 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 **apm-owa-session-irule**.
4. In the **Definition** section, copy and paste the following iRule:

```
when ACCESS_ACL_ALLOWED {
   set apm_mrhsession [HTTP::cookie value "MRHSession"]
   if { [table lookup $apm_mrhsession] == "EXCHANGE_LOGOUT" } {
      ACCESS::session remove
      table delete $apm_mrhsession
   }
   unset apm_mrhsession
}

when HTTP_REQUEST {
   if {[string tolower [HTTP::uri]] contains "owa" } {
      if {[string tolower [HTTP::uri]] contains "logoff.aspx" } {
         ACCESS::session remove
         HTTP::redirect "https://[HTTP::host]/owa"
      } else {
      set isset 0
      set request_uri [HTTP::uri]
      if { [HTTP::uri] contains "UA=0" } {
         set mrhsession [HTTP::cookie value "MRHSession"]
         set isset 1
         }
      }
   }
}

when HTTP_RESPONSE {
   if { $isset == 1 } {
      if { $mrhsession != "" && [HTTP::status] == 440 } {
         table set $mrhsession "EXCHANGE_LOGOUT"
         unset mrhsession
         unset request_uri
         return
      }
         unset isset
         unset mrhsession
   }
   unset request_uri
}
```

5. Click **Finished**.

Modifying the virtual server to use the APM persistence iRule

The next task is to modify the virtual server created by the template to use the iRules you just created.

To modify the virtual server

1. On the Main tab, expand **Local Traffic** and then click **Virtual Servers**.
2. From the **Virtual Server** list, click the name of the single virtual server for the Client Access HTTP-based services on port 443 created by the template. By default, this is **my_Exchange_2010__single_https_virtual_server**.
3. On the Menu bar, click **Resources**.
4. In the iRules section, click **Manage**.
5. From the **Enabled** box, select the Persistence iRule create by the template and then click the Remove (**>>**) button. By default this is **my_Exchange_2010__single_Persist_irule**.
6. From the **Available** list, click the name of the persistence iRule you created and then click the Add (**<<**) button to move it to the Enabled box.
7. From the **Available** list, click the name of the iRule you created to terminate inactive sessions, and then click the Add (**<<**) button to move it to the Enabled box.
8. Click the **Finished** button.

After creating the iRule, continue with the configuration table on page 45.

Scenario 4: Configuring the BIG-IP APM (SSL Bridging)

In this scenario, you are using a single BIG-IP configured with both APM and LTM modules. 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.

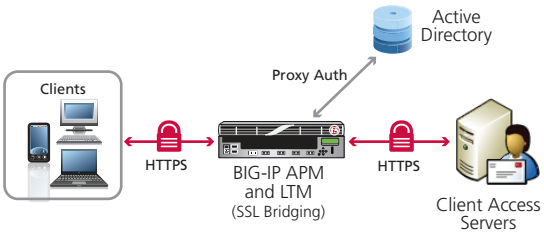

Figure 2.4: *BIG-IP APM with SSL Bridging configuration example*

For scenario 4, you must complete the following procedures:

- *Deleting the redirect virtual server, on this page*
- *Creating new HTTPS monitors, on this page*
- *Creating the iRules, on this page*
- *BIG-IP Edge Gateway and APM Configuration on page 45*
- *Configuring the BIG-IP APM for Scenarios 3 and 4 on page 48*

Deleting the redirect virtual server

For this scenario, we recommend you delete the redirect virtual server on the BIG-IP LTM because it is not used. Use the procedure in Scenario 2, *Deleting the redirect virtual server on the BIG-IP LTM on page 38* to delete the port 80 virtual server.

Creating new HTTPS monitors

Because the template creates HTTP monitors for the Client Access HTTP-based services, for this scenario, you must create new HTTPS monitors for those services.

To create the new monitors, use the procedure from Scenario 2, *Creating new HTTPS monitors on the BIG-IP LTM on page 38*. Be sure to also use the procedure to modify the pools to use the new monitors.

Creating the iRules

The next task is to create a new persistence iRule, and the iRule to terminate inactive APM sessions, and then modify the virtual server to use the iRules. To create the persistence iRule, use the procedure *Creating the persistence iRule when using BIG-IP APM on page 40.*
To create the inactive sessions iRule, use the procedure *Creating the iRule to terminate inactive sessions on page 42*

Modifying the virtual server to use the APM persistence iRule

The next task is to modify the virtual server created by the template to use the iRule you just created.

To modify the virtual server to use the APM persistence iRule, use the procedure *Modifying the virtual server to use the APM persistence iRule on page 42*.

After completing these changes, continue with the configuration table on page 45.

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

[1] 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.
   a. Click the **Logon Page** option button, and then the **Add Item** button at the bottom.
   b. Configure the Logon Page properties as applicable, and then click **Save**. In our example, we leave the defaults.

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.

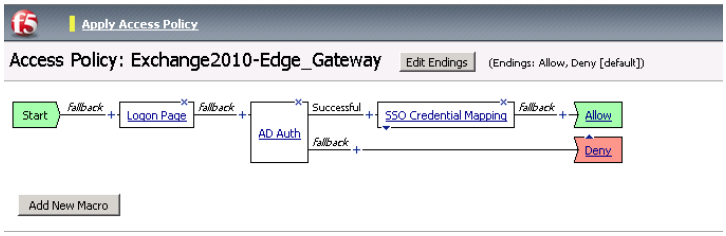

Creating the iRule that chooses the SSO Configuration

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

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             WEBSSO::select $static::OWA_FORM_BASE_SSO_CFG_NAME
10  }
11      unset req_uri
12  }
```

4. Click the **Finished** button.

Configuring the Edge Gateway for Scenarios 1 and 2

If you are using the BIG-IP Edge Gateway for scenarios 1 (SSL offload) or 2 (SSL Bridging), use the following table to configure the Edge Gateway.

BIG-IP Object	Non-default settings/Notes	
Health Monitor *(Main tab-->Local Traffic-->Monitors)*	***Name***	Type a unique name. We use **edge-exchange-monitor**.
	Type	TCP
Pool (Main tab-->Local Traffic-->Pools)	***Name***	Type a unique name. We use **edge-exchange-pool**.
	Health Monitor	Name of the health monitor you created above
	New Member Address	The BIG-IP LTM virtual server IP address(es) for the Client Access HTTP-based services created by the template
	Service Port	**80** (**443** if not offloading SSL)
Profiles (Main tab-->Local Traffic-->Profiles)	*OneConnect (Profiles-->Other-->OneConnect)*	
	Name	Type a unique name. We use **edge-exchange-oneconnect**.
	Parent Profile	**oneconnect**; all other settings at default levels.
	HTTP (Profiles-->Services->HTTP)	
	Name	Type a unique name. We use **edge-exchange-http**.
	Parent Profile	**http**
	Redirect Rewrite	For scenario 1 (offload), select **All** from the list.. For scenario 2 (SSL bridging) **do not** modify Redirect Rewrite
	Client SSL (Profiles-->SSL-->Client SSL)	
	Name	Type a unique name. We use **edge-exchange-clientssl**.
	Parent Profile	**clientssl**
	Certificate and Key	Select the SSL certificate and key you imported.
	Server SSL (Profiles-->SSL-->Server SSL) For scenario 2: SSL Bridging only	
	Name	Type a unique name. We use **edge-exchange-serverssl**.
	Parent Profile	**serverssl**
	Certificate and Key	Use the default certificate and key.
Edge Gateway Exchange virtual server (Main tab-->Local Traffic-->Virtual Servers)	***Name***	Type a unique name. We use **edge-exchange-virtual**.
	Destination Address	The IP address clients use to access Microsoft Exchange. Your Exchange FQDN resolves to this IP address.
	Service Port	**443**
	OneConnect profile	Select the OneConnect profile you created above.
	HTTP Profile	Select the HTTP profile you created above
	SSL Profile (Client)	Select the Client SSL profile you created above.
	SSL Profile (Server)	Select the Server SSL profile you created above (only for Scenario 2, SSL Bridging).
	Access Profile	Select the Access Profile you created above
	iRules	Enable the built-in **_sys_APM_ExchangeSupport_OA_BasicAuth** iRule, Enable the inactive session iRule you created Enable the SSO iRule you created (**select_SSO_irule** in our example) Enable the APM session ID irule you created (**edge-gateway-irule** in our example). If deploying ActiveSync, enable the built in **_sys_APM_activesync** iRule.
	Default Pool	Select the Pool you created above

Important → *After configuring the Edge Gateway as described in the table above, you may need to modify your FQDN in DNS to point to the Edge Gateway Exchange Virtual Server (as described in the table) and not the BIG-IP LTM virtual server.*

Configuring the BIG-IP APM for Scenarios 3 and 4

If you are using the BIG-IP APM for scenarios 3 (SSL offload) or 4 (SSL Bridging), you need to modify the single virtual server for Exchange 2010 HTTP-based services to use the BIG-IP APM objects you created. If you are configuring the APM for scenario 4 (SSL Bridging), you must first create a Server SSL profile.

Configuring a Server SSL profile for scenario 4 (SSL Bridging)

For SSL bridging, you must create a Server SSL profile. Use the following table

BIG-IP Object	Non-default settings/Notes	
Server SSL profile *(Main tab-->Local Traffic-->Profiles-->SSL -->Server SSL)*	***Name***	Type a unique name. We use **exchange-serverssl**.
	Parent Profile	**serverssl**
	Certificate and key	Use the default certificate and key

Modifying the LTM virtual server to use the APM objects for Scenarios 3 and 4

The final task for scenarios 3 and 4 is to modify the BIG-IP LTM virtual server created by the template to use the APM objects, and iRules you created.

To modify the BIG-IP LTM virtual server for the Exchange 2010 HTTP-based services

1. On the Main tab, expand **Local Traffic** and then click **Virtual Servers**.
2. From the **Virtual Server** list, click the name of the single virtual server for the Client Access HTTP-based services on port 443 created by the template. By default, this is **my_Exchange_2010__single_https_virtual_server**.
3. If necessary, from the **Configuration** list, select **Advanced**.
4. *For scenario 4 (SSL Bridging)* ***only***:
   From the **SSL Profile (Server)** list, select the Server SSL profile you just created.
5. From the **Access Profile** list, select the name of the Access profile you created earlier in this section.
6. Click the **Update** button.
7. On the Menu bar, click **Resources**.
8. In the iRules section, click **Manage**.
9. From the **Enabled** box, select the Persistence iRule create by the template and then click the Remove (**>>**) button. By default this is **my_Exchange_2010__single_Persist_irule**.
10. From the **Available** box, select the following iRules and then click the Add (**<<**) button to move them to the **Enabled** box.
    a. The persistence iRule you created in *Creating the persistence iRule when using BIG-IP APM on page 40*.
    b. The iRule that selects the SSO configuration you created in *Creating the iRule that chooses the SSO Configuration on page 46*.
    c. The built-in iRule **_sys_APM_ExchangeSupport_OA_BasicAuth**.
    d. If deploying ActiveSync, the built-in iRule **_sys_APM_activesync**.
11. Click the **Finished** button.

This completes the configuration for Edge Gateway and BIG-IP APM for all scenarios.

Configuring the FirePass controller for Client Access servers

This section of the Deployment Guide shows you how to configure the F5 FirePass controller for secure remote access to the Microsoft Exchange Server 2010 Client Access Servers.

F5's FirePass® controller is the industry leading SSL VPN solution that enables organizations of any size to provide ubiquitous secure access for employees, partners and customers to applications such as Microsoft Exchange Server 2010, while significantly lowering support costs associated with legacy client-based VPN solutions.

Prerequisites and configuration notes

The following are prerequisites for this section:

- The FirePass controller should be running version 6.0 or later.
- All of the configuration procedures in this section are performed on the FirePass controller.
- This configuration uses previously defined Active Directory groups to provide authentication and simple user maintenance. For information on how to configure Active Directory groups, consult the appropriate documentation.
- This Deployment Guide is written to the scenario outlined in the following section. It is meant as a template; modify the configuration as necessary for your deployment.

Configuration scenario

For the scenario used in this Deployment Guide, the Microsoft Exchange deployment, along with an Active Directory instance, resides behind a BIG-IP system. A group on the FirePass controller is given three access methods for reading Microsoft Exchange/Outlook Web Access email:

- Through an Outlook Web Access Portal Favorite on the FirePass device.
- Through the Network Access adapter, with a locally installed Microsoft Outlook client.
- Through the Mobile Email feature, which provides lightweight, pure HTML access to Exchange mailboxes using IMAP/POP3 and SMTP.

The table on the following page describes how to configure the FirePass controller to allow secure remote access to the Exchange device(s), using Active Directory for authentication. In our deployment, the FirePass device and the Exchange deployment use a common Active Directory Domain Controller.

Configuration table for FirePass

The table on the following page shows the necessary objects to configure on the FirePass controller, and any non-default settings. For specific instructions on how to configure individual objects, see the online help or product documentation.

Configuration table for FirePass

FirePass Object	Description/Notes
Resource Group	Create a **Resource Group**. - Name = Type a unique name. All other settings are optional.
Master Group	Create a **Master Group**. - *Users in Group* = **External** - *Authentication method* = **Active Directory** - *Copy Settings from* = **Do not copy** After clicking Create: - Click *Resource Groups* tab and add resource group you created above. - Click *Authentication* tab. Configure with your Active Directory settings. - Click *Select Domain Group*, and select the appropriate Active Directory Domain group.
Auto-logon/Single sign-on	From **Portal Access**, open the **Master Group Settings** for the group you created. Configure the following: - Check the box to *Limit Web Applications Access to Intranet Favorites only...* - Check the box in *Auto-login to Basic and NTLM auth protected sites...*
Outlook Web App[1]	Click **Users**-->**Groups**-->**Resource Groups**-->*Name of your resource group.* - In the *Portal Access* column, click **Edit**. - Create a new *Web Application Favorite*, using the **Microsoft Outlook Web Access** *Web Application.* Enter the URL used to access OWA.
Mobile Email[2]	Click **Portal Access-->Mobile E-Mail** - Check the **Enable corporate mail account box**. - Type a name for the account. - In the *Mail Server* box, type the name or IP of the Exchange server. - From *Type* list, select **IMAP**. - In the *IMAP Folders* box, type the appropriate folders. - From the *Login Information* list, select the appropriate setting. - In the *Outgoing Mail Server* box, type the name or IP of the outgoing mail server. - All other settings are optional.
Network Access[3]	Click **Network Access-->Global Settings** - Create IP address Pool in *Add new IP address Pool* section: Name, IP address[3], Mask. Click **Apply these rules now**. Click **Network Access-->Resources** - In Client Settings tab ~ Connection Name: Type a name. This is what end users see in their Favorite list. ~ **Client for Microsoft Networks** = checked. Split Tunneling, File/Printer Sharing, and GZIP are optional. Click **Update**. ~ In Configure IP Address Assignment section, **Assign IP address dynamically using IP address pool** = checked ~ In Select IP Address Pool section, select the IP Address Pool you created above, and click **Update**.
Endpoint Security	Configure Endpoint Security as applicable for your implementation. We recommend creating a pre-logon sequence that includes at least an anti-virus check. For information on configuring Endpoint security and pre-logon sequences, see the Firepass documentation or online help.

[1] *For organizations who want an added layer of security for their Outlook Web App deployment, want to require antivirus or other pre-logon checks, or do not want to make Outlook Web Access directly accessible from the Internet, the FirePass can be configured to render Outlook Web Access inside the FirePass user window.*

[2] *As an alternative (or in addition to) using Outlook Web App, you can use the FirePass controller's Mobile Email feature as a lightweight and extremely secure way of viewing Microsoft Exchange email.*

[3] *Using Network Access requires one internal IP address for each concurrent user, so make sure the Network IP address can handle all possible concurrent users (for example, 10.10.101.0 creates enough addresses for 254 users.*
To prevent routing problems, ensure the Network address pool does not contain the FirePass device's IP address.

Chapter 3

Deploying F5 and Microsoft Exchange Server 2010 Edge Transport Servers

This chapter gives you guidance on configuring F5 products for deployment with the Edge Transport Server component of Exchange Server 2010.

In Exchange 2010, the Edge Transport server role is usually deployed in your organization's perimeter network on standalone servers. Designed to minimize the attack surface, the Edge Transport server handles all Internet-facing mail flow, which provides Simple Mail Transfer Protocol (SMTP) relay and smart host services for the Exchange organization. Edge Transport servers also include anti-spam and antivirus features, which provide services to block viruses and spam, or unsolicited commercial e-mail, at the network perimeter.

For more information on the Edge Transport Server role, see *technet.microsoft.com/en-us/library/dd351247.aspx*

For more information on Microsoft Exchange Server 2010, see *www.microsoft.com/exchange/default.mspx*.

For more information on F5 products and features, see *http://www.f5.com/products/*.

Prerequisites and configuration notes

The following are prerequisites and configuration notes for this chapter:

- This chapter contains procedures on configuring multiple F5 products and/or modules. To perform certain procedures, you must own the appropriate product or licensed the relevant module. These sections are clearly marked.
- All of the configuration procedures in this chapter are performed on F5 devices. For information on how to deploy or configure Microsoft Exchange Server 2010, consult the appropriate Microsoft documentation.

Configuration example

As Edge Transport Servers are most often located on or near the perimeter of an organization's networks, it is possible to deploy Edge Transport servers in more than one datacenter. Any or all of those Edge Transport servers may be involved in relaying mail.

In the following deployment, the BIG-IP LTM system provides local traffic management and uses SMTP health monitors to check the availability of the Edge Transport servers. We also use the Message Security Module (MSM) to provide the first line of defense in the fight against SPAM. MSM can eliminate up to 70% of unwanted email before letting the Edge Transport servers handle the rest.

We also enable the GTM module in two data centers, set up active monitoring of the status of Local Traffic Manager virtual servers that are in front of Edge Transport server pools, establish a DNS record for the mail service, and build policies which direct incoming email appropriately.

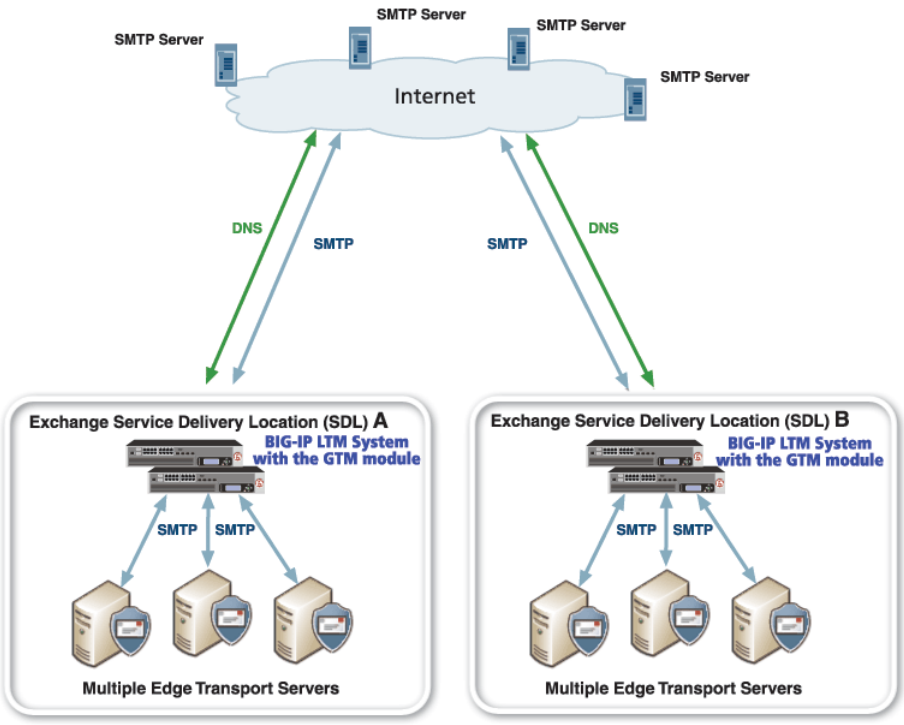

Figure 3.1: *BIG-IP and Edge Transport Server logical configuration example*

Configuration table for BIG-IP LTM: Edge Transport Servers

The following table shows the non-default settings on BIG-IP LTM objects for the Edge Transport Servers. For specific instructions on how to configure individual objects, see the online help or product documentation.

Virtual Server	Pool	Health monitor	TCP profile
Service Port : **25** (or select **SMTP** from the list) Add Pool and TCP profile	LB Method: We recommend *Least Connections (node)* Service Port : **25** (or select **SMTP** from the list)	**SMTP**: Base SMTP monitor. *Domain*: Enter appropriate domain. *Interval*: We recommend 30 *Timeout*: We recommend 91	Base TCP profile (all other settings optional)

Using the Message Security Module for Edge Transport Servers

The Message Security Module (MSM) identifies and blocks unwanted emails at the edge of your network. You configure MSM to block known and malicious spam senders, and keep them from filling your network with unwanted email. Blocking unwanted email at the edge of your network minimizes the resource load on your network and associated devices like Exchange Server 2010.

MSM includes a real-time subscription to McAfee® TrustedSource®, and email filtering capabilities for the BIG-IP system. TrustedSource is an industry-leading system for evaluating the safety of email sources, and for scoring the reputation of the IP addresses from which email originates.

We recommend that you use MSM as a spam volume-control solution in addition to using the existing, content-based, email filtering solutions that are already installed on your network, like those provided with Exchange 2010 Edge Transport servers. This combination provides more complete protection for your network than either solution alone.

For more information on the Message Security Module, see the documentation available on Ask F5.

MSM prerequisites and configuration notes

The following are prerequisites and configuration notes specific to the Message Security Module:

- You must have purchased the Message Security Module. For more information about purchasing the Message Security Module, contact your sales representative.
- MSM is available with BIG-IP LTM version 9.4 and later.
- We assume you have already installed and licensed the Message Security Module. For more information on installing and licensing MSM, see the MSM documentation available on Ask F5.
- You must have command line access to the Root directory of the BIG-IP system. This means that you must be assigned the Administrator.

Configuring the MSM to manage email traffic to Edge Transport servers

The BIG-IP Message Security Module installation creates a data group named **MSM_config**, and adds the following three variables and default attributes to the data group:

- trusted_pool:good_mail
- suspect_pool:maybe_mail
- quarantine_pool:quarantine_mail

These variables correspond to the IP address reputation scores that TrustedSource assigns to the sources requesting connection to your network (for more information on reputation scores, see the MSM documentation appropriate for your BIG-IP version, such as *http://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/msm1_1_deploy.html*).

The default value for each variable is the name of a pool of mail servers to which MSM directs a specified kind of traffic, as shown in the following table.

Variable in MSM_config	Default Value
trusted_pool	**good_mail** This is the name of the pool of mail servers to which MSM load balances mail from trusted sources.
suspect_pool	**maybe_mail** This is the name of the pool of mail servers to which MSM load balances mail from suspect sources. That is, mail that you want your existing email filtering systems to scan.
quarantine_pool	**quarantine_mail** This is the name of the pool of mail servers to which MSM load balances mail that you want to quarantine on your network for possible manual analysis.

Creating the pools

You can create the three pools described in preceding table, or you can use existing pools to manage your email traffic. For Edge Transport servers, we recommend you create two new pools for the suspect and quarantine mail; trusted mail is sent to the existing Edge Transport pool you created previously (requires modifying the MSM_config data group, as shown later in this section). If you want to use other existing pools, you need to modify the names of variables in the MSM_ config data group for these pools as well.

Create two new pools, named **maybe_mail** and **quarantine_mail**. Add the appropriate mail servers to the pool. All other settings are optional. For specific instructions on how to create a pool, see the online help or the product documentation.

Modifying the names of variables in the MSM_config data group

The next task is to modify the MSM_config data group to send trusted mail to the Exchange 2010 Edge Transport server pool we already created. If the other pools you created have different names than the pool names in strings, you must modify these string records as well.

To modify the MSM_config data group

1. On the Main tab, expand **Local Traffic**, and then click **iRules**.
2. On the menu bar, click **Data Group List**.
3. In the **Name** column, click **MSM_config**.
4. In the Records area, modify the string records that represent the load balancing pools that handle the email on your system.
   a. In the **String Records** list, select the **trusted_pool:good_mail** string, and then click the **Edit** button.
   b. Change **good_mail** to the name of the pool you created in the *Configuration table for BIG-IP LTM: Edge Transport Servers on page 25*. In our example, we change it to be **trusted_pool:exch_et_pool**.
   c. Click the **Add** button.
   d. Repeat steps a-c to modify any of the other default names.

5. Optional: You can also modify the strings that represent the threshold values. Follow the same substeps as above. For more information on the threshold values, see the MSM or TrustedSource documentation.
6. Click **Finished**.
7. After updating the data group, you must force MSM to re-initialize the class data. To do this:
   a. Open an SSH client and log in to the BIG-IP system as an administrator.
   b. Run the following command from the command line:
      ```
      # MSM_init
      ```
      This loads the MSM data class and initializes the new values.

Modifying the virtual server to use the MSM

The final task is to modify the virtual server you created for the Edge Transport Servers to reference the iRule that the BIG-IP Message Security Module installation process creates.

To modify the virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
2. From the **Virtual Server** list, click the name of the virtual server you created in *Configuration table for BIG-IP LTM: Edge Transport Servers on page 25*. In our example, we select **exch_et_virtual**.
3. From the **Configuration** list, select **Advanced**.
4. From the **Statistics Profile** list, and select **MSM_reputation**. (This is the Statistics profile that the MSM installation process created.)
5. Click the **Update** button.
6. On the menu bar, click **Resources**.
7. In the iRules section, click the **Manage** button.
8. From the **iRules Available** list, select **MSM_reputation**, and click the Add (<<) button to move the iRule to the Enabled list. (This is the iRule that the MSM installation process created.)
9. Click the **Finished** button

Configuring the BIG-IP GTM for Edge Transport Servers

The Edge Transport role for Microsoft Exchange 2010 provides inbound and outbound SMTP connectivity between an Exchange organization and other mail services, including all other Internet email users. More information on the Edge Transport server role may be found at *http://technet.microsoft.com/en-us/library/dd351247.aspx*.

Most often located on or near the perimeter of an organization's networks, it is possible to deploy Edge Transport servers in more than one datacenter. Any or all of those Edge Transport servers may be involved in relaying mail.

Traditional methods of providing high availability to public-facing SMTP mail relays involve using a combination of simple round-robin DNS and multiple MX (mail exchange) DNS records that statically list two or more delivery locations, with fixed priority levels. Those methods do not provide true load balancing, do not permit dynamic redirection based on performance, and make it difficult to perform maintenance or cope with localized outages in predictable and controllable ways.

Using F5's Global Traffic Manager (GTM) allows mail administrators to define policies which take into account real-time availability and performance of all Edge Transport servers, plan and easily initiate local maintenance outages without disrupting service, and remain highly available even in the event of a disaster.

To use the BIG-IP GTM for Microsoft Exchange 2010, you should have two or more data centers in your deployment that will host globally load balanced Edge Transport servers.

Configuration table for BIG-IP GTM: Edge Transport Servers

The table on the following page shows the non-default settings on BIG-IP GTM objects for the Edge Transport Servers. For specific instructions on how to configure individual objects, see the online help or product documentation.

Important →

You need a unique self IP address for each redundant pair of BIG-IP LTM devices in this configuration. If you have multiple pairs of BIG-IP LTMs you need a unique self IP for each one. The IP address you choose, and the VLAN to which you assign it, must be accessible by any clients that will be performing DNS queries against the GTM. It may be a private IP address if a Network Address Translation (NAT) device, such as a BIG-IP LTM, a firewall, or a router, is providing a public address and forwarding DNS traffic to the listener.

GTM Object	Description/Notes
Listener	Create a **Listener** for each BIG-IP (or redundant pair) in this configuration. The *Destination* is the Self IP address on the BIG-IP LTM system (in the Important Note on the previous page). *VLAN Traffic* = **All VLANs**.
Data Center	One GTM **Data Center** for each location that will host globally load balanced Edge Transport servers. *Name* is required, all other settings are optional.
Health monitor	*Type* = **SMTP** monitor (recommended) All settings are optional.
Server	One GTM **Server** for each data center *Product* = Either **BIG-IP System (Single)** or **BIG-IP System (Redundant)** *Address* = BIG-IP LTM Self IP Address. *Data Center* = Data Center created above *Health Monitor* = Monitor created above. *Virtual Server Discovery* = **Enabled (No Delete)** (recommended not required) All other settings optional.
Pool	One GTM **Pool** for each LTM virtual server. *Health Monitor* = Monitor created above. *Load Balancing Method* = Preferred: **Global Availability**, Alternate: **Round Robin**, Fallback: **Return to DNS** (recommended) *Member List-->Virtual Server* list: Add all BIG-IP LTM virtual servers containing the Edge Transport servers.
Wide IP	*Load Balancing* = **Global Availability** (recommended) *Pool List* = Add all GTM Pools created in previous section All other settings optional.
Zonerunner (or other DNS system)	Add the Wide IP created above as a MX record in DNS (if not using GTM, consult the appropriate documentation or consult your DNS administrator). If using ZoneRunner on the GTM, create a new **Resource Record**. *Name* = Add a name. Make sure the domain for which you are creating an MX record is shown, and note that it must end with a period *TTL* = Type a time to live. We use 500 *Type* = MX *Preference* = 10. *Mail Server* = Wide IP you created in the previous section

For more information or specific instructions on creating GTM objects, see the Online Help or GTM documentation.

Chapter 4

Deploying BIG-IP WOM with Exchange 2010 DAG and Hub Transport Servers

In this chapter, we provide configuration parameters for configuring the BIG-IP WAN Optimization Module (WOM) for deployment with the Mailbox Server Database Availability Group (DAG) and Hub Transport features of Exchange Server 2010.

BIG-IP WOM for Database Availability Groups

A DAG can contain up to sixteen member servers; you can configure continuous replication of mailbox databases copies between any of those members. Because DAG members can be located across a WAN in alternate data centers, wide-area bandwidth and latency can affect the replication rate. If replication does not keep up with data generation, an outage or failover to an alternate DAG member can result in data loss.

For more information on the WAN Optimization Module, see *http://www.f5.com/products/big-ip/wan-optimization-module.html*

By default, DAG replication is configured for network encryption and compression between members on separate subnets. By disabling that default setting and configuring WOM to perform those functions, you gain the following benefits:

- By offloading CPU and memory overhead associated with compression and encryption from your Exchange Mailbox servers, you allow them to spend more resources on Mailbox role functionality.
- By combining BIG-IP's TCP Express network optimization with WOM compression and symmetric deduplication, you speed up encrypted mailbox database replication between DAG members, reducing or eliminating the potential downtime or data loss caused by replication that is not up-to-date at the time of an outage.
- You reduce the total amount of data transferred over the WAN connection, enabling other operations to proceed more efficiently over the same connection or reducing overall bandwidth needs.

More information about DAGs, including requirements and Exchange Server configuration, can be found at: *technet.microsoft.com/en-us/library/dd979799%28EXCHG.140%29.aspx*

BIG-IP WOM for Hub Transport

Servers with the Hub Transport role installed are responsible for routing email between two or more Exchange sites in an organization. These messages can form a large part of the traffic on an organization's WAN. By using BIG-IP WOM to optimize the WAN traffic, messages are routed more reliably and with less impact on your organizations WAN links.

More information about the Hub Transport Server role can be found at *http://technet.microsoft.com/en-us/library/aa998616.aspx*

Supported Topologies for DAG

BIG-IP WOM can accommodate several different network topologies for Exchange Server 2010 DAG replication networks. The core configuration steps on the BIG-IP WOM remain the same in all cases; however, specific configuration details required during configuration, such as routes, Self IPs, local endpoints, and advertised local subnets, will vary according to your specific network design.

Note

BIG-IP supports high-availability configurations; e.g. a pair of BIG-IP devices at either or both sides of the WAN. For simplicity in the following diagrams, we show only one BIG-IP in each location, each of which has only one IP address per subnet. In a high-availability scenario, you would still perform all of the steps listed in this chapter, but you would also configure individual Self IP addresses for each BIG-IP. The iSession Self IP and the Self IP configured as a gateway in the following diagrams and procedures would be configured as a "Floating IP". See the BIG-IP product documentation for more information on configuring high availability pairs.

Supported network topologies include:

BIG-IP WOM as the gateway on the DAG Networks; single subnet for your WAN

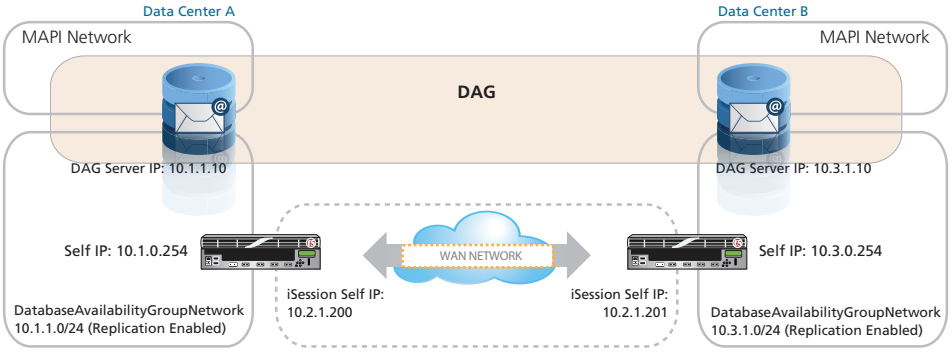

Figure 4.1: *WOM as the Gateway; single subnet*

BIG-IP WOM as the gateway on the DAG Networks; separate subnets for either side of WAN

Note

The examples for DAG in the rest of this chapter refer to this diagram

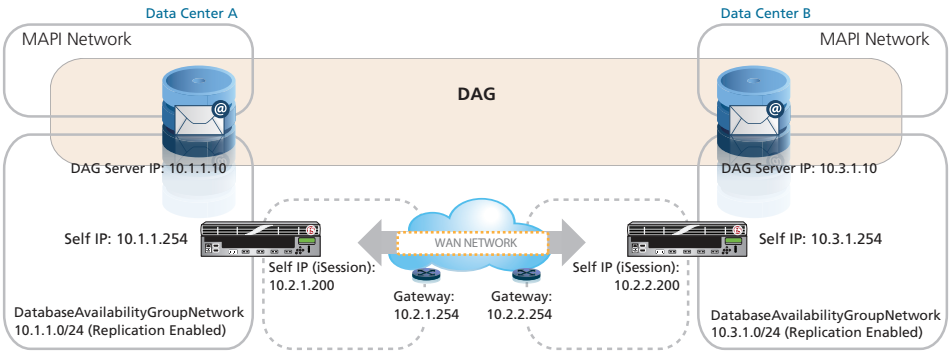

Figure 4.2: *WOM as the Gateway, separate subnets*

A router as the gateway on the DAG Networks, with BIG-IP WOM on separate subnet

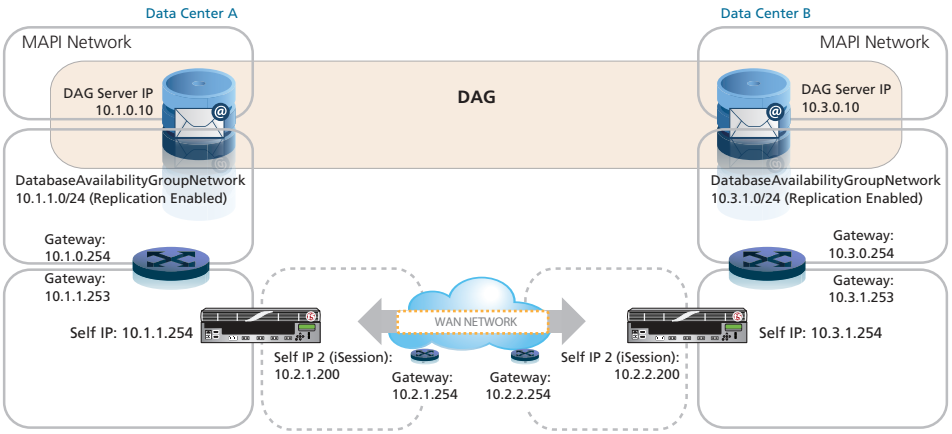

Figure 4.3: *Router as the Gateway, WOM on separate subnets*

Supported topologies for Hub Transport

BIG-IP WOM can optimize inter-site Hub server mail transport over any supported topology. As with optimizing DAG, the core configuration steps on the BIG-IP WOM remain the same in all cases; however, specific configuration details required during configuration, such as routes, Self IPs, local endpoints, and advertised local subnets, will vary according to your specific network design. For simplicity, we show a single common topology.

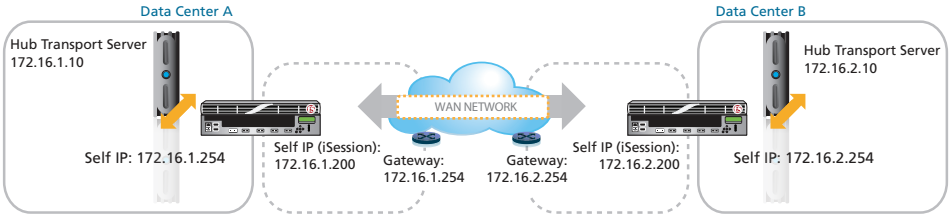

Figure 4.4: *WOM as the gateway, single subnet*

Configuring the BIG-IP WOM

The following diagram visually represents the order in which the WOM devices and Exchange Server 2010 must be configured whether you are optimizing DAG or Hub traffic. You should have access to the BIG-IP WOM devices and Exchange servers in both data centers.

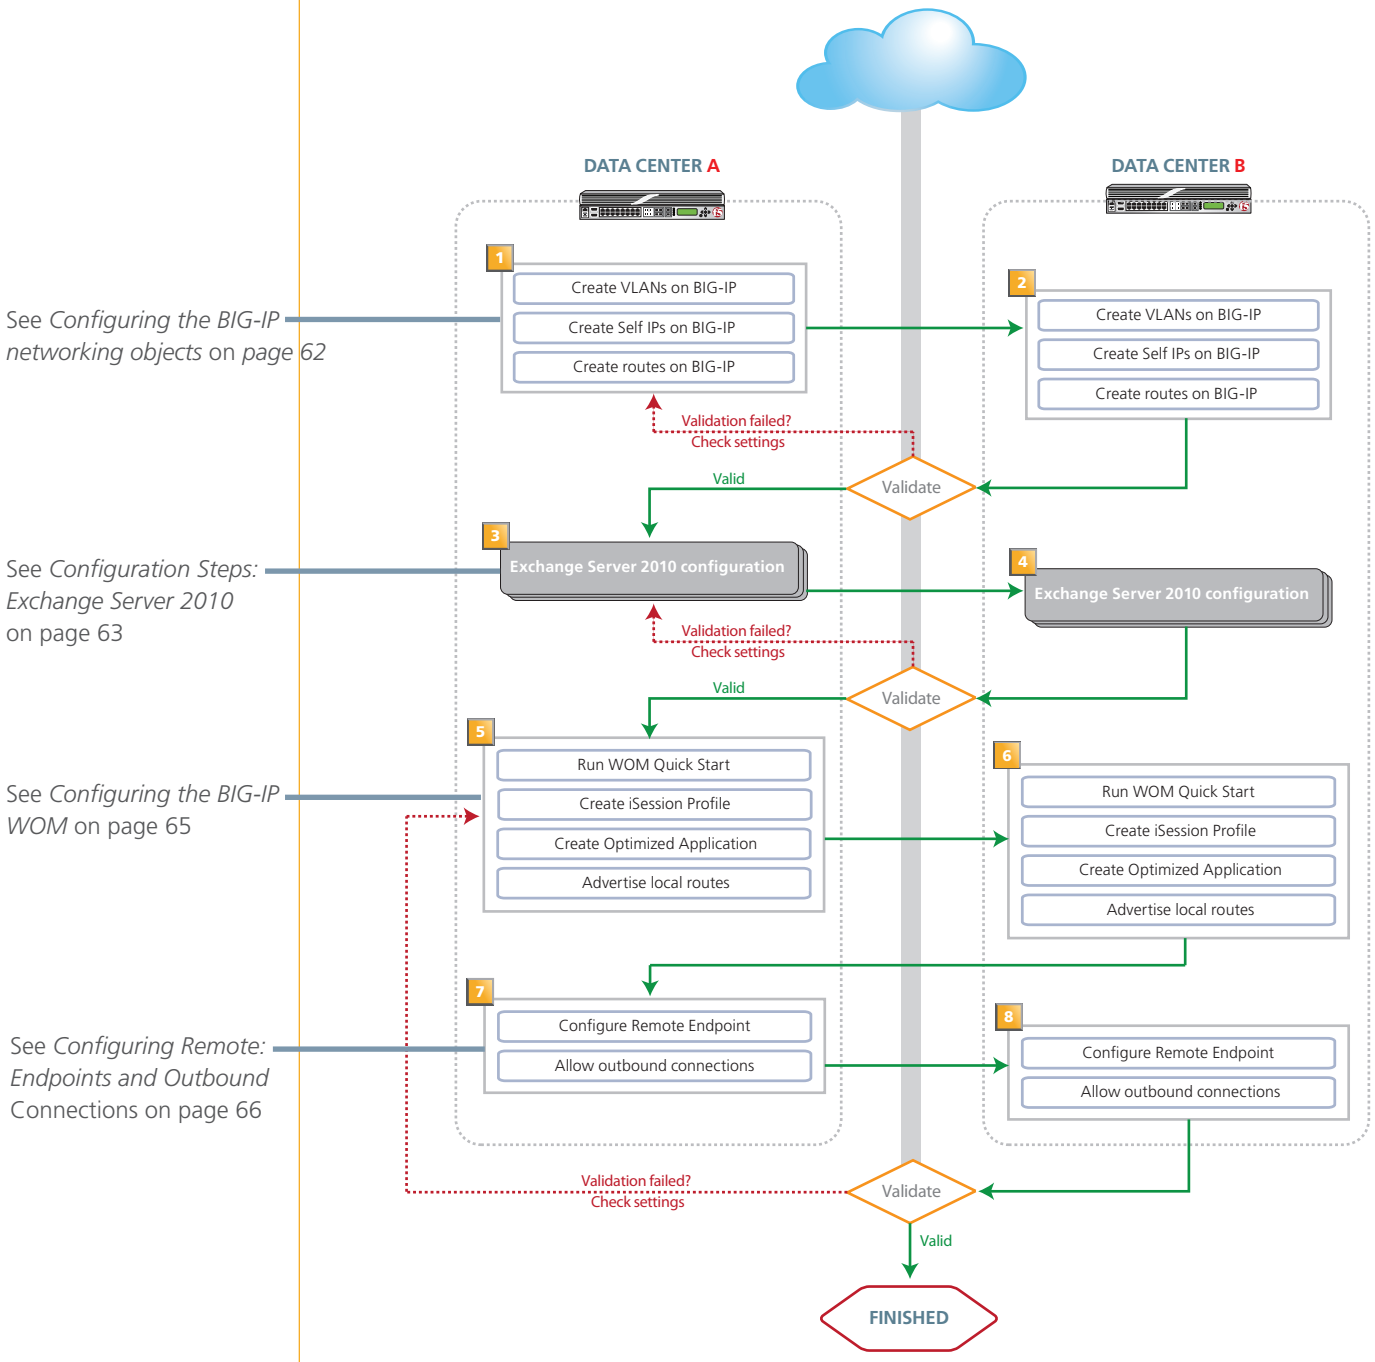

See *Configuring the BIG-IP networking objects* on *page 62*

See *Configuration Steps: Exchange Server 2010* on page 63

See *Configuring the BIG-IP WOM* on page 65

See *Configuring Remote: Endpoints and Outbound* Connections on page 66

Figure 4.5: *Configuration flow diagram*

Configuring the WOM networking objects

In this section, we create the networking objects on the BIG-IP systems in both data centers. Use the following table for guidance on each object. For specific information on how to configure individual objects, see the online help or product documentation.

WOM Object	Description/Notes
VLANs	WOM requires at least two VLANs: a LAN VLAN and a WAN VLAN. Give each VLAN a descriptive name. Note: We find using VLAN tags makes management easier. However, tagging is not mandatory if your configuration can support individual interfaces instead of VLANs.
Self IPs	Assign an otherwise-unused static IP address that resides on the VLAN you created. The LAN-side Self IP is used as a gateway on that network. The WAN-side Self IP is used for the WOM iSession; (we also refer to this as the Local Endpoint Self IP). For the iSession VLAN ***only***: From the **Port Lockdown** list, select **Allow None**. For all other VLANs, Port Lockdown should be set to **Allow Default**. In a high availability configuration, you assign individual Self IP addresses to each BIG-IP in each subnet first, and then assign the additional Self IP addresses listed described here, but designate them as Floating IPs. Only the iSession Self IP is configured with **Port Lockdown: None**.
Routes	Create a route on BIG-IP in the primary data center to route to the BIG-IP in the secondary data center. You also need a route for the remote network where application services reside. For example, in Figure 2 the BIG-IP in data center A would need a route to the address 10.2.2.254 (the self IP of the remote BIG-IP) and the 10.3.1.0/24 network (the remote DAG network), both using the gateway 10.2.1.254. The BIG IP in data center B needs corresponding routes back to the BIG-IP and the DAG network in data center A.

Repeat the WOM networking objects in the secondary data center

Next, repeat the VLAN, Self IP and Route configuration on the BIG-IP WOM in the secondary data center.

ⓘ **Checkpoint**

Use the checkpoints to ensure the configuration thus far is working properly.

Checkpoint

At this point, you should be able to ping the local endpoint, the router, and the BIG-IP system in the secondary data center. The following requires command line or SSH access to the BIG-IP system. Consult the product documentation for instructions on how to use SSH.
The following examples refer to Figure 4 on page 31.

- Log into the BIG-IP in the primary data center from the command line.
- Use the ping command to check the following. You should receive successful responses from each:
  - Local endpoint - In our example, we use **`ping 10.1.0.254`**
  - The router - In our example, we use **`ping 10.2.1.254`**
  - The BIG-IP in the secondary data center - In our example, we use **`ping 10.2.2.200`**.
- Repeat this procedure in the other data center using the appropriate IP addresses.

If you do not receive successful responses, check the IP addresses and VLAN configuration.

Configuration Steps: Exchange Server 2010

In the following sections, we describe how to configure Exchange 2010 properties to enable acceleration by WOM. We are providing configuration information for Exchange Server 2010 only as a general guide; for detailed instructions you should refer to the Microsoft resources that we link to in some of the steps, or related official product documentation. Where relevant, we indicate information that you will need in subsequent steps to configure BIG-IP WOM.

Configuring WOM and DAG

Tip →

Where we show Exchange Management Shell commands, in most cases you can also use the Exchange Management Console (EMC) to perform the same tasks from a graphical interface.

Before proceeding, you should obtain all the information required to complete the following worksheet. This information is used in both the DAG and BIG-IP WOM configuration.

	Data Center A - Primary	Data Center B - Cloud	Notes
DAG Network and Subnet Mask			
Self IP of BIG-IP on DAG Network			Applicable only if BIG-IP is directly on DAG network
WAN network and Subnet Mask			
WAN network upstream gateway IP address			

Our example, based on Figure 2, looks like the following:

	Data Center A - Primary	Data Center B - Cloud	Notes
DAG Network and Subnet Mask	10.1.1.0 255.255.255.0	10.3.1.0 255.255.255.0	
Self IP of BIG-IP on DAG Network	10.1.1.254	10.3.1.254	Applicable only if BIG-IP is directly on DAG network
WAN network and Subnet Mask	10.2.1.0/24	10.2.2.0/24	
WAN network upstream gateway IP address	10.2.1.254	10.2.2.254	

In the case of a BIG-IP high availability pair, the Self IP referred to in the table would be the Floating IP address.

Optimizing an existing DAG

You should follow the instructions in this section if you have an existing DAG and are implementing WOM to accelerate and secure database replication. If you have not yet configured a DAG, refer to the section Creating a New DAG.

Because compression and encryption will be handled by the BIG-IP devices, you must disable Exchange Server's native DAG Network Compression and DAG Network Encryption; these parameters are set on a per-DAG level, not for individual servers. Use the following command syntax within the Exchange Management shell (all on one line):

```
Set-DatabaseAvailabilityGroup –Identity <DAG name> -NetworkCompression Disabled –
NetworkEncryption Disabled
```

See *http://technet.microsoft.com/en-us/library/dd298065.aspx* for full instructions on the **Set-DatabaseAvailabilityGroup** cmdlet.

Creating a new DAG

If you have not yet created a DAG, use the following steps.

1. Install and configure Mailbox servers in each datacenter according to Exchange Server 2010 documentation and best practices. At a minimum, each of the Mailbox servers must have two (2) network interfaces in separate subnets; one will be used for MAPI, and one for the DAG. The IP address and subnet of the DAG networks will be needed for WOM configuration; see the table on the previous page.
2. Designate an otherwise-unused IP address on the MAPI network to assign to the DAG in each data center.
3. Create a DAG; assign it a name and the MAPI-network IP addresses you designated in step 2, above. To do so, in the Shell, type the following all on one line:

   ```
   New-DatabaseAvailabilityGroup –Name <DAGname> -DatbaseAvailabilityGroupIPAddress
   <Datacenter A DAG address>,<Datacenter B DAG address>
   ```

   You may also need to assign a witness server. Complete instructions for creating a DAG can be found at *http://technet.microsoft.com/en-us/library/dd351172.aspx*.
4. Configure DAG networks according to the requirements listed at *http://technet.microsoft.com/en-us/library/dd638104.aspx*. Note that the maximum round-trip latency for a DAG network cannot exceed 500ms. .
5. Add each of your Mailbox servers, in both datacenters, to the DAG. In the Shell, use the following command syntax (all on one line):

   ```
   Add-DatabaseAvailabilityGroupServer –Identity <DAGname> -MailboxServer
   <servername>
   ```

   You can find instructions for adding members to the DAG at *http://technet.microsoft.com/en-us/library/dd298049.aspx*
6. Because compression and encryption will be handled by the BIG-IP devices, you must ensure that DAG Network Compression and DAG Network Encryption are disabled; these parameters are set on a per-DAG level, not for individual servers. Use this command within the Exchange Management shell:

   ```
   Set-DatabaseAvailabilityGroup –Identity <DAG name> -NetworkCompression Disabled
   –NetworkEncryption Disabled
   ```

   See *http://technet.microsoft.com/en-us/library/dd298065.aspx* for full instructions on the **Set-DatabaseAvailabilityGroup** cmdlet.

7. DAG members on different subnets may have more than one IP address used by the host servers. You must deploy your BIG-IP devices within your network to optimize traffic between your DAG networks rather than between your MAPI or other networks. You can set the DAG replication via either the shell, using the **Set-DatabaseAvailabilityGroupNetwork** cmdlet and the **–ReplicationEnabled** parameter, or via the EMC. You can find more information on DAG network replication at *http://technet.microsoft.com/en-us/library/dd297927.aspx*.

The remaining configuration procedures detailed in this chapter are performed on the F5 devices. For further information on how to deploy or configure Microsoft Exchange Server 2010, consult the appropriate Microsoft documentation.

Configuring WOM and Inter-site Hub Transport

To configure Hub mail transport to use WOM for WAN optimization, you must first disable the default inter-site TLS (encryption) settings on Hub servers. Follow the instructions in either of these two documents:
"Disabling TLS Between Active Directory Sites to Support Wan Optimization"
http://technet.microsoft.com/en-us/library/ee633456.aspx

"Suppress Anonymous TLS Connections"
http://technet.microsoft.com/en-us/library/dd876856.aspx

Before proceeding, you should obtain all the information required to complete the following worksheet.

	Data Center A - Primary	Data Center B - Cloud	Notes
Hub Server Network and Subnet Mask			
Self IP of BIG-IP on the Hub Server Network			Applicable only if BIG-IP is directly on Hub Server network
WAN network and Subnet Mask			
WAN network upstream gateway IP address			

In our example, the table looks like the following:

	Data Center A - Primary	Data Center B - Cloud	Notes
Hub Server Network and Subnet Mask	172.16.1.0 255.255.255.0	172.16.2.0 255.255.255.0	
Self IP of BIG-IP on the Hub Server Network	172.16.1.254	172.16.2.254	Applicable only if BIG-IP is directly on Hub Server network
WAN network and Subnet Mask	172.16.1.0/24	172.16.2.0/24	
WAN network upstream gateway IP address	172.16.1.254	172.16.2.254	

In the case of a BIG-IP high availability pair, the Self IP referred to in the table would be the Floating IP address.

Configuring the BIG-IP WAN Optimization settings

The next task is to configure the BIG-IP WOM. Use the following table for guidance on each object. For specific information on how to configure individual objects, see the online help or product documentation.

WOM Object	Description/Notes
WOM Quickstart	**WAN Self IP Address:** Self IP address for the Local Endpoint **Discovery:** Enabled **LAN VLAN**: Select the VLAN that corresponds to the local DatabaseAvailabilityGroupNetwork **WAN VLAN**: Select the VLAN that corresponds to the WAN side of your replication network. **Authentication and Encryption** section: Leave the defaults. **Create Optimized Applications** section: Do **NOT** check the box for Microsoft Exchange (MAPI).
iSession Profile	Give the profile a unique name. If you are encrypting DAG replication traffic (strongly recommended), Enable **Application Data Encryption**. Leave all other settings at the default levels.
Optimized Application	Give the Optimized Application a unique name. **Port: 64327** (the default DAG replication port, if you changed the DAG properties to use another port, type that port. **Enabled LAN VLANs**: Choose the local DAG network VLAN **iSession Profile**: Choose the iSession profile you created above.
Advertised Routes	**Address**: This is the IP address that specifies the local subnet you have configured in your DAG for replication. **Netmask**: The corresponding subnet mask.

Repeat the WOM configuration in the secondary data center

Next, repeat the WOM quickstart, iSession profile, Optimized Application and Advertised Routes configuration on the BIG-IP WOM in the secondary data center.

Configuring remote endpoints and outbound connections

The final tasks in the WOM configuration are configuring the remote endpoints and confirming that outbound connections are allowed.

WOM Object	Description/Notes
Remote Endpoint	**IP address**: The Self IP address for the BIG-IP WOM module in the secondary data center (this is what you configured as the iSession Self IP on the remote WOM).
Confirming Outbound Connections are allowed	From the Remote Endpoints section of the GUI, click the IP address of the Endpoint you created above. - Make sure there is a green circle in the left column. If it is red, there is a connectivity problem; recheck connectivity between data centers. - In the *Outbound iSession to WAN* section, make sure there is a check in the **Outbound Connections** box. If there is not, check the box and then click **Update**.

Repeat the remote endpoint configuration in the secondary data center
Next, repeat the remote endpoint and outbound connection configuration on the BIG-IP WOM in the secondary data center.

ⓘ **Checkpoint**
Use the checkpoints to ensure the configuration thus far is working properly.

Checkpoint: Testing the configuration

At this checkpoint, we make sure that BIG-IP WOM connectivity between the primary and secondary data center is enabled. As this is a critical point in this configuration, we use two different procedures to make sure the WOM tunnel is properly configured.

To test WOM connectivity

1. From the Main tab of the BIG-IP configuration utility, expand **WAN Optimization**, and then click **Remote Endpoints**.
2. From the list of Remote Endpoints, make sure the Status Indicator is green for the endpoint in the secondary data center.
3. If the status indicator is a color other than green, on the Main tab under WAN Optimization, and then click Diagnostics. Run through all of the troubleshooting diagnostics.

Use the following follow this procedure to ensure that the WOM tunnel endpoints are up and running properly. For the procedure you will need SSH access to the BIG-IP.

To verify the WOM tunnel

1. Using an SSH client, like Putty, establish a connection to each BIG-IP.
2. After logging in, at the command prompt, type **`b endpoint remote show all`**
   You should see an output similar to the following, however your host name and IP addresses will be different. Make sure you see the tunnel state as **ready, ready**.

```
b endpoint remote show all

ENDPOINT REMOTE 20.20.20.20
| HOSTNAME PRIMARYDC.example.com
| MGMT ADDR 10.1.102.61 VERSION 10.2.0
| UUID c1f3:68d6:f697:6834:108:5668:1e16:3fce
| enable STATE ready (incoming, outgoing)=(ready, ready)
| BEHIND NAT disable
| CONFIG STATUS “none”
| DEDUP CACHE 62380 REFRESH (count) = (0)
| ALLOW ROUTING enable
+-> ENDPOINT REMOTE 20.20.20.20 ROUTE 20.20.20.0/24
| | INCLUDE enable LABEL West
```

3. SSH to the second BIG-IP and verify the tunnel status shows ready/ready.

➲ **Note:** *Only proceed with configuration after the status of the tunnel shows* **ready/ready**.

Appendix A: Manual configuration tables

We recommend using template to configure the BIG-IP system for the Client Access server role. This table contains the BIG-IP configuration objects in this deployment and any non-default settings for advanced users. See Chapter 2 for Edge Gateway and APM configuration.

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

CAS Role/BIG-IP object	Non-default settings/Notes		
Health Monitors	***Outlook Web App***	**HTTP** parent (use **HTTPS** for SSL Bridging) see recommended health monitor configuration in Step 4 on *page 8*)	
	Outlook Anywhere	**HTTP** parent (use **HTTPS** for SSL Bridging) see recommended health monitor configuration in Step 4 on *page 10)*	
	ActiveSync	**HTTP** parent (use **HTTPS** for SSL Bridging) see recommended health monitor configuration in Step 4 on *page 12)*	
	Autodiscover	**HTTP** parent (use **HTTPS** for SSL Bridging) see recommended health monitor configuration in Step 4 on *page 14*)	
Pools ***(repeat for each CAS role)***	***Health monitor***	Add the appropriate health monitor for the Client Access role you created above	
	Slow Ramp Time	**300** (must select Advanced from the Configuration menu for this option to appear)	
	Load Balancing Method	**Least Connections (member)** recommended	
	Address	IP Address of Client Access server running Outlook Web App	
	Service Port	**80** (repeat Address and Port for all members) `Important`*: Create a pool for each Client Access Server role*	
iRules	Append[1] (*page 75), Persistence iRule* (*page 75).* `Important`*: The Append iRule should be listed first when configuring the virtual server*		
Profiles	***HTTP***	Parent Profile	If using WebAccelerator: **http-acceleration**. If not using WebAccelerator: **http-wan-optimized-compression-caching**
		Redirect Rewrite	**All**
		Compression-->Content List-->Include List *(Add each entry to the Content Type box and click Include)*	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|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)
		RAM Cache-->Content List-->Exclude List	/owa/ev.owa
			uglobal.js
			oab.xml
	TCP WAN[1]	Parent Profile	**tcp-wan-optimized**
	TCP LAN[1]	Parent Profile	**tcp-lan-optimized**
	Client SSL	Parent Profile	**clientssl**
		Certificate/Key	Select the Certificate and Key you imported
	Server SSL[2]	Parent Profile	**serverssl**
	Persistence[3]	Persistence Type	**Cookie**
	OneConnect	Parent Profile	**oneconnect**
		Source Mask	**255.255.255.255**
	NTLM	Parent Profile	**ntlm**
Virtual Servers	***Port 443***	Destination Address	IP address for the virtual server (Service Port 443)
		Profiles	Add select each of the profiles you created above from the appropriate list
		SNAT Pool	**Automap**[4]
		iRules	Append, Persistence (the Append iRule must be listed first)
		Default Pool	Do ***not*** select a default pool for this virtual
	Port 80 (optional, for redirect purposes only)	Destination Address	IP address for the virtual server (Service Port 80)
		Profiles	HTTP profile only
		iRule	_sys_https_redirect

[1] *The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent*
[2] *Server SSL profile is only necessary if configuring SSL Bridging.*
[3] ***Important****: See Required: Modify the Cookie persistence profile timeout value on page 24 for an important modification to this profile.*
[4] *See Using a SNAT Pool if you expect more than 6,000 users per Client Access server on page 26*

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

Outlook Web App configuration table

BIG-IP object	Non-default settings/Notes		
Health Monitor	***Type***	**HTTP** (use **HTTPS** for SSL Bridging) See optional, recommended health monitor configuration in Step 4 on *page 8*)	
Pool	***Health monitor***	Add health monitor above	
	Slow Ramp Time	**300** (must select Advanced from the Configuration menu for this option to appear)	
	Load Balancing Method	**Least Connections (member)** recommended	
	Address	IP Address of Client Access server running Outlook Web App	
	Service Port	**80** (repeat Address and Port for all members)	
iRules	Append[2] *(page 75) and* the built-in iRule:_**sys_http_redirect**[1] if offloading SSL		
Profiles	***HTTP***	Parent Profile	If using WebAccelerator: **http-acceleration**. If not using WebAccelerator: **http-wan-optimized-compression-caching**
		Redirect Rewrite	**All**
		Content List-->Include List *(Add each entry to the Content Type box and then click Include)*	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|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)
		Content List-->Exclude List	/owa/ev.owa
			uglobal.js
			oab.xml
	TCP WAN[1]	Parent Profile	**tcp-wan-optimized**
	TCP LAN[1]	Parent Profile	**tcp-lan-optimized**
	Client SSL	Parent Profile	**clientssl**
		Certificate/Key	Select the Certificate and Key you imported
	Server SSL[2]	Parent Profile	**serverssl**
	Persistence[3]	Persistence Type	**Cookie**
	OneConnect	Parent Profile	**oneconnect**
		Source Mask	**255.255.255.255**
	NTLM	Parent Profile	**ntlm**
Virtual Servers	***Port 443***	Destination Address	IP address for the virtual server (Service Port 443)
		Profiles	Add select each of the profiles you created above from the appropriate list
		SNAT Pool	**Automap**[4]
		iRules	Append, Persistence (the Append iRule must be listed first)
		Default Pool	Select the pool you created for Outlook Web App above
	Port 80 (optional, for redirect purposes only)	Destination Address	IP address for the virtual server (Service Port 80)
		Profiles	HTTP profile only
		iRule	_sys_https_redirect

[1] *The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent*
[2] *Server SSL profile is only necessary if configuring SSL Bridging.*
[3] ***Important**: See Required: Modify the Cookie persistence profile timeout value on page 24 for an important modification to this profile.*
[4] *See Using a SNAT Pool if you expect more than 6,000 users per Client Access server on page 26*

Outlook Anywhere manual configuration table (for separate virtual server configuration)

BIG-IP object	Non-default settings/Notes		
Health Monitor	***Type***	**HTTP** (use **HTTPS** for SSL Bridging) see recommended health monitor configuration in Step 4 on *page 10*)	
Pool	***Health monitor***	Add health monitor above	
	Slow Ramp Time[1]	**300**	
	Load Balancing Method	**Least Connections (member)** recommended	
	Address	IP Address of Client Access server running Outlook Anywhere	
	Service Port	**80** (repeat Address and Port for all members)	
iRules	**OA Persist** *(page 78)* which is associated with the persistence profile below *and* the built-in iRule:**_sys_http_redirect**[1] if offloading SSL		
Profiles	***HTTP***	Parent Profile	**http**
		Redirect Rewrite	**All**
	TCP WAN[2]	Parent Profile	**tcp-wan-optimized**
	TCP LAN[2]	Parent Profile	**tcp-lan-optimized**
	Client SSL	Parent Profile	**clientssl**
		Certificate/Key	Select the Certificate and Key you imported
	Persistence	Persistence Type	**Universal**
		iRule	Select the OA Persist iRule you created above
	OneConnect	Parent Profile	**oneconnect**
		Source Mask	**255.255.255.255**
	NTLM	Parent Profile	**ntlm**
Virtual Servers	***Port 443***	Destination Address	IP address for the virtual server (Service Port 443)
		Profiles	Add select each of the profiles you created above from the appropriate list
		SNAT Pool	Automap[4]
		Default Pool	Select the pool you created for Outlook Anywhere above
	Port 80 (optional, for redirect purposes only)	Destination Address	IP address for the virtual server (Service Port 80)
		Profiles	HTTP profile only
		iRule	_sys_https_redirect

Active Sync manual configuration table (for separate virtual server configuration)

BIG-IP object	Non-default settings/Notes		
Health Monitor	***Type***	**HTTP** (use **HTTPS** for SSL Bridging) see recommended health monitor configuration in Step 4 on *page 12*)	
Pool	***Health monitor***	Add health monitor above	
	Slow Ramp Time[1]	**300**	
	Load Balancing Method	**Least Connections (member)** recommended	
	Address	IP Address of Client Access server running ActiveSync	
	Service Port	**80** (repeat Address and Port for all members)	
Profiles	***HTTP***	Parent Profile	**http**
	TCP WAN[2]	Parent Profile	**tcp-wan-optimized**
	TCP LAN[2]	Parent Profile	**tcp-lan-optimized**
	Client SSL	Parent Profile	**clientssl**
		Certificate/Key	Select the Certificate and Key you imported
	Server SSL[3]	Parent Profile	**serverssl**

[1] *You must select Advanced from the Configuration list for this option to appear*
[2] *The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent*
[3] *Server SSL profile is only necessary if configuring SSL Bridging.*
[4] *See Using a SNAT Pool if you expect more than 6,000 users per Client Access server on page 26*

Active Sync manual configuration table for separate virtual server configuration (continued)

BIG-IP object	Non-default settings/Notes		
Virtual Servers	***Port 443***	Destination Address	IP address for the virtual server (Service Port 443)
		Profiles	Add select each of the profiles you created above from the appropriate list
		SNAT Pool	**Automap**[3]
		Default Pool	Select the pool you created for ActiveSync above
	Port 80 (optional, for redirect purposes only)	Destination Address	IP address for the virtual server (Service Port 80)
		Profiles	HTTP profile only
		iRule	_sys_https_redirect

Autodiscover manual configuration table (for separate virtual server configuration)

BIG-IP object	Non-default settings/Notes		
Health Monitor	***Type***	**HTTP** (use **HTTPS** for SSL Bridging) see recommended health monitor configuration in Step 4 on *page 14*)	
Pool	***Health monitor***	Add health monitor above	
	Slow Ramp Time[1]	**300**	
	Load Balancing Method	**Least Connections (member)** recommended	
	Address	IP Address of Client Access server running Autodiscover	
	Service Port	**80** (repeat Address and Port for all members)	
Profiles	***HTTP***	Parent Profile	**http**
	TCP WAN[2]	Parent Profile	**tcp-wan-optimized**
	TCP LAN[2]	Parent Profile	**tcp-lan-optimized**
	Client SSL	Parent Profile	**clientssl**
		Certificate/Key	Select the Certificate and Key you imported
	Server SSL[3]	Parent Profile	**serverssl**
Virtual Servers	***Port 443***	Destination Address	IP address for the virtual server (Service Port 443)
		Profiles	Add select each of the profiles you created above from the appropriate list
		SNAT Pool	**Automap**[4]
		Default Pool	Select the pool you created for Autodiscover above
	Port 80 (optional, for redirect purposes only)	Destination Address	IP address for the virtual server (Service Port 80)
		Profiles	HTTP profile only
		iRule	_sys_https_redirect

[1] *You must select Advanced from the Configuration list for this option to appear*
[2] *The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent*
[3] *Server SSL profile is only necessary if configuring SSL Bridging.*
[4] *See Using a SNAT Pool if you expect more than 6,000 users per Client Access server on page 26*

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[1] dynamic port range manual configuration table

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

[1] *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.*

[2] *You must select Advanced from the Configuration list for this option to appear*

[3] *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.*

[4] *See Using a SNAT Pool if you expect more than 6,000 users per Client Access server on page 26*

RPC Client Access[1] static ports configuration table

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

[1] *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.*

[2] *You must select Advanced from the Configuration list for this option to appear*

[3] *The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent*

[4] *See Using a SNAT Pool if you expect more than 6,000 users per Client Access server on page 26*

POP3 manual configuration table

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

IMAP4 manual configuration table

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

[1] *You must select Advanced from the Configuration list for this option to appear*
[2] *The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent*
[3] *Server SSL profile is only necessary if configuring SSL Bridging.*
[4] *See Using a SNAT Pool if you expect more than 6,000 users per Client Access server on page 26*

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.

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.

OWA Redirect iRule (formerly referred to as the Append iRule)

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

ActiveSync persist iRule

If you are deploying ActiveSync on a BIG-IP behind a NAT or other address aggregating device, use this iRule to ensure even distribution of client connections.

```
1  when HTTP_REQUEST {
2      if { [HTTP::header exists "Authorization"] } {
3          persist uie [HTTP::header "Authorization"] 7200
4      } else {
5          persist source_addr
6      }
7  }
```

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 to match the names of the pools in your configuration.*

If you selected LAN when asked from where clients are primarily connecting, you MUST remove or comment out the CACHE::disable line where specified in the iRule
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/BIG-IPv10-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***.

Critical *You must change the pool names to match the names of the pools in your configuration*

```
## iRule to select pool and persistence method when all Exchange Client
## Access HTTP-based services are accessed through the same BIG-IP virtual
## server. This iRule will use an HTTP header inserted by a BIG-IP Edge
## Gateway for persistence (if that header is present); otherwise it will
## set persistence according to traditional methods.

## CHANGE ALL POOL NAMES TO MATCH THOSE IN YOUR ENVIRONMENT.

when HTTP_REQUEST {

    ## Offline Address Book and Autodiscover do not require persistence.
    switch -glob -- [string tolower [HTTP::path]] {

        "/microsoft-server-activesync" {
            ## ActiveSync.
            if { [HTTP::header exists "APM_session"] } {
                persist uie [HTTP::header "APM_session"] 7200
            } elseif { [HTTP::header exists "Authorization"] } {
                set as_key [sha256 [HTTP::header "Authorization"]]
                persist uie $as_key 7200
            } else {
                persist source_addr
            }
            pool as_pool_name
            COMPRESS::disable
            ## If you selected LAN when asked from where clients are primarily
            ## connecting, you MUST remove or comment out the CACHE::disable line
            CACHE::disable
            return
        }

        "/owa*" {
            ## Outlook Web Access
            if { [HTTP::header exists "APM_session"] } {
                persist uie [HTTP::header "APM_session"] 7200
            } else {
                persist cookie insert
            }
            pool owa_pool_name
            return
        }

        "/ecp*" {
            ## Exchange Control Panel.
            if { [HTTP::header exists "APM_session"] } {
                persist uie [HTTP::header "APM_session"] 7200
            } else {
                persist cookie insert
            }
            pool owa_pool_name
            return
        }

        "/ews*" {
            ## Exchange Web Services.
            if { [HTTP::header exists "APM_session"] } {
                persist uie [HTTP::header "APM_session"] 7200
            } else {
                persist source_addr
            }
            pool oa_pool_name
            COMPRESS::disable
            ## If you selected LAN when asked from where clients are primarily
            ## connecting, you MUST remove or comment out the CACHE::disable line
            CACHE::disable
            return
        }
```

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

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

```
        "/oab*" {
            ## Offline Address Book.
            pool oa_pool_name
            return
        }

        "/rpc/rpcproxy.dll" {
        ## Outlook Anywhere.
            if { [HTTP::header exists "APM_session"] } {
                persist uie [HTTP::header "APM_session"] 7200
                } elseif { [string tolower [HTTP::header "Authorization"]] starts_with "basic" } {
                set oa_key [sha256 [HTTP::header "Authorization"]]
                persist uie $oa_key 7200
                } else {
                persist source_addr
                }

            pool oa_pool_name
            COMPRESS::disable
            ## If you selected LAN when asked from where clients are primarily
            ## connecting, you MUST remove or comment out the CACHE::disable line
            CACHE::disable
            return
        }

        "/autodiscover*" {
            ## Autodiscover.
            pool ad_pool_name
            return
        }

        default {
            ## This final section takes all traffic that has not otherwise
            ## been accounted for and sends it to the pool for Outlook Web App
            if { [HTTP::header exists "APM_session"] }  {
                persist uie [HTTP::header "APM_session"] 7200
            } else {
                persist source_addr
            }
            pool owa_pool_name
        }
    }
}

when HTTP_RESPONSE {
    if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate"} {
        ONECONNECT::reuse disable
        ONECONNECT::detach disable
        ## this command disables NTLM conn pool for connections where OneConnect has been disabled
        NTLM::disable
    }
    ## this command rechunks encoded responses
    if {[HTTP::header exists "Transfer-Encoding"]} {
        HTTP::payload rechunk
    }
}
```

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.

```
when HTTP_REQUEST {
   switch -glob -- [string tolower [HTTP::path]] {
      "/ews*" {
      ## Exchange Web Services.
         if { [HTTP::header exists "APM_session"] } {
            persist uie [HTTP::header "APM_session"] 7200
            } else {
            persist source_addr
            }
      }

      "/rpc/rpcproxy.dll" {
      ## Outlook Anywhere.
         if { [HTTP::header exists "APM_session"] } {
            persist uie [HTTP::header "APM_session"] 7200
            } elseif { [string tolower [HTTP::header "Authorization"]] starts_with "basic" } {
            persist uie [HTTP::header "Authorization"] 7200
            } else {
            persist source_addr
            }
      }
   }
}

when HTTP_RESPONSE {
    if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate"} {
        ONECONNECT::reuse disable
        ONECONNECT::detach disable
        ## disables NTLM conn pool for connections where OneConnect has been disabled
        NTLM::disable
    }
    ## this command rechunks encoded responses
    if {[HTTP::header exists "Transfer-Encoding"]} {
        HTTP::payload rechunk
    }
}
```

Appendix B: Technical Notes

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

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 Exchanger 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.

Although the BIG-IP GUI cannot create SAN certificates, and will not display the Subject Alternative Name values of imported certificates, use of SAN certificates is otherwise supported.

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	

Appendix C: Using X-Forwarded-For to log the client IP address in IIS 7.0 and 7.5

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. By configuring an HTTP profile on the BIG-IP to insert an X-Forwarded-For header, 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.

Your BIG-IP LTM HTTP profile must have X-Forwarded-For enabled. See *Required: Modifying the HTTP profile on page 23*, or the HTTP profile section of the manual configuration table.

Deploying the Custom Logging role service

The next 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**.
5. Click the **Next** button.
6. On the Confirmation page, click **Install**.
7. After the role service has successfully installed, you can 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

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/Joe/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 in the Add Logging Field box, and then click the **OK** button in the Edit Logging Fields box.
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.

This completes the configuration.

Deployment Guide Revision History

Version	Description	Date
1.0	Original deployment guide for Exchange Server 2010.	N/A
1.1	Added support for BIG-IP v10.1. Replaced previous guidance with BIG-IP WOM for DAG and Hub Transport.	N/A
1.2	- Added support for the Autodiscover service. - Removed the OneConnect profile from RPC Client Access, Outlook Anywhere, and POP3/IMAP4. - Updated the following iRules: Append, Outlook Anywhere persistence, and the combined virtual server persistence; added Appendix A with a full version of the combined iRule.	N/A
1.3	Corrected an error in the Outlook Anywhere Persistence iRule and same error in the combined virtual server persistence iRule.	N/A
1.4	- Added guidance that when repeating the OWA virtual server procedure for other Exchange services, you should not add the Appending iRule to the Outlook Anywhere, ActiveSync or Autodiscover virtual servers. - Modified the Configuration table for Outlook Anywhere to clarify the Persist iRule is associated with the Universal Persistence profile, and not the virtual server.	N/A
1.5	Corrected an error in the single virtual server persistence iRule on page 1-41. This error was not present in the commented version of the iRule in Appendix A or in the iRule produced by the Exchange 2010 Application template.	N/A
2.0	- Complete document rewrite and reformat. - Removed many configuration procedures in favor of Application Template guidance and Configuration tables.	N/A
2.1	- Added configuration to disable OneConnect for Kerberos-authenticated connections (Autdiscover) on *page 26*. - Changed the Exchange Server settings table, Autodiscover Authentication row, Your Setting column, to **Negotiate**.	N/A
2.2	Corrected the optional monitor Send String for OWA *page 8* (step 4b) to include a space between the 0 and HTTP. With no space, the monitor would fail even though the service was up.	N/A
3.0	- Substantially updated the configuration for Edge Gateway and APM - Updated manual configuration tables in the appendix. - Modified timeout value for RPC Client Access persistence from 3600 to 7200 seconds. - Added Appendix B: Technical Notes with detailed descriptions of Slow Ramp Time, SAN SSL certificates and Outlook Client configuration. - Added updated iRules for persistence.	N/A
3.1	Modified the persistence iRules (for both LTM and Edge/APM) with the following changes: - Changed "/Microsoft-Server-ActiveSync" to "/microsoft-server-activesync" (note case) - Changed "/xml/autodiscover.aspx" to " /autodiscover*" - In the "/rpc/rpcproxy.dll"section, added **`string tolower`** to the first else statement. - Changed "MSRPC" to "msrpc" - Changed "*Microsoft Office" to "*microsoft office"	N/A
3.2	Added the persistence iRule to a ZIP file to avoid any potential white space errors.	N/A
3.3	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. Updated the SNAT Pool iRule on *page 26*.	N/A
3.4	Modified the persistence iRule if using a single virtual server (lines 102 and 103 on page 74) to change "header" to "cookie".	N/A
3.5	- 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. - Reformatted the cover page and chapter title pages.	01-17-2012
3.6	Added guidance for enabling the built in **_sys_APM_ExchangeSupport_OA_BasicAuth** iRule to the BIG-IP LTM virtual server in *Configuring the BIG-IP APM for Scenarios 3 and 4 on page 48*.	01-23-2012
3.7	- Added guidance for enabling the built-in **_sys_APM_activesync** iRule when deploying APM and ActiveSync. - Clarified SNAT Pool iRule guidance on page 24. The iRule uses IP addresses added to the SNAT Pool, and not new addresses.	03-07-2012
3.8	- 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 Log Definition to include the X-Forwarded-For header. Moved this section to Appendix C. - 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.	03-13-2012
3.9	*Manual configuration only*: Modified the HTTP 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 **Keep Accept Encoding** box. You should not enable Keep Accept Encoding.	04-19-2012

Version	Description	Date
4.0	- Updated the persistence iRules with multiple fixes. - Updated the External Monitor script file for Autodiscover, and added a new file if using SSL Bridging. - Added the following items to *Modifying the template configuration on page 22:* • Modifying the OneConnect profile to include a Source Mask of 255.255.255.255. • Modifying the Cookie Persistence timeout value. • Modifying the Persistence Profile for RPC Client Access to include enabling Match Across Virtual Servers. • Added instructions for adding the RPC Client Access persistence profile to either the combined virtual server or single Outlook Anywhere virtual server as a fallback persistence profile. - Updated the manual configuration tables to include SSL Bridging option and to reflect all the changes above. - Added support for BIG-IP v10.2.4. We strongly recommend using v10.2.4 and later.	05-22-2012
4.1	Updated *Optional: Creating an EAV monitor for Autodiscover on page 28* to include the correct directory for the monitor if using BIG-IP version 10.1 or later in the 10.x branch.	06-21-2012
4.2	Added a new iRule for Edge Gateway and APM deployments only that terminates inactive sessions. See *Creating the iRule to terminate inactive sessions on page 42.*	06-26-2012
4.3	- Corrected the preparation worksheet on page 5 to reference the correct Client Access Service for ActiveSync and Autodiscover - Added to the procedure *Required: Modifying the HTTP profile on page 23* to include an addition step if you selected clients are coming over a LAN when configuring the template for OWA or the single virtual server. - Added the procedure *Required: Modifying the Append iRule for OWA on page 25* to the Modifying the Template Configuration section. - Added the procedure *Required: Adding the ActiveSync persist iRule if using separate virtual servers on page 27* to the Modifying the Template Configuration section. - Added the procedure *Modifying the IIS authentication token timeout value on page 28* to the Modifying the Template Configuration section.	10-03-2012
4.4	Updated the document to include guidance for modifying the iRules if you chose LAN as the location where users are primarily connecting. In this case, you must comment out the CACHE::disable line in the iRule. Added comments to the iRules with instructions. Related to the same issue, added guidance in *Required: Modifying the HTTP profile on page 23* to select a the http_optimized_compression_caching parent profile if you selected WAN.	11-16-2012
4.5	Added two new URIs to the Exclude list (/owa/ev.owa and oab.xml) to step 8 of *Required: Modifying the HTTP profile on page 23* in the modifying the template configuration section. Added the same URIs to the Exclude list in the manual configuration tables.	05-13-2013
4.6	In the manual configuration table for RPC Client Access with static ports, added an Idle Timeout value of 7200. There is an existing section in the Modifying the Template configuration for this change, but it was missing from the manual configuration table for static ports.	07-08-2013
4.7	- Corrected the name of the parent HTTP profile in Step 3a of *Required: Modifying the HTTP profile on page 23,* and clarified the guidance in Step 7. - Corrected the parent HTTP profile in the manual configuration tables for Outlook Web App on page 68 and 69.	07-09-2013
4.8	Modified the optional, advanced health monitor Send and Receive Strings for ActiveSync on page 12.	07-11-2013
4.9	- Added an additional entry to the Modifying the Template Configuration section concerning iOS users experiencing certificate errors after deploying the template for ActiveSync. See *Replacing the SSL profile if users are having trouble with iOS and ActiveSync on page 30.* - Updated the applicable iRules to use SHA256 (replacing CRC32), as it is more secure.	01-22-2014
4.9.1	Removed a comment line about TMM sending gratuitous ARPs during failover from the SNAT Pool iRule on *page 26*	02-20-2014
4.9.2	Added support for Exchange 2010 SP3	06-09-2014

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com