

Microsoft Exchange Client Access Servers

Welcome to the F5 and Microsoft® Exchange® 2010 and 2013 Client Access Server deployment guide. Use this document for guidance on configuring the BIG-IP system version 11 and later to provide additional security, performance and availability for Exchange Server 2010 and Exchange Server 2013 Client Access Servers.

When configured according to the instructions in this guide, whether using an iApp template or manually, the BIG-IP system performs as a reverse proxy for Exchange CAS servers, and also performs functions such as load balancing, compression, encryption, caching, and pre-authentication.

Why F5?

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

- The BIG-IP LTM can balance load and ensure high-availability across multiple Client Access servers using a variety of load-balancing methods and priority rules.
- Terminating HTTPS connections at the BIG-IP LTM reduces CPU and memory load on Client Access Servers, and simplifies TLS/SSL certificate management for Exchange 2010 and Exchange 2013 SP1 and later.
- The BIG-IP Access Policy Manager (APM), F5's high-performance access and security solution, can provide preauthentication, single sign-on, and secure remote access to Exchange HTTP-based Client Access services.
- The BIG-IP Advanced Firewall Manager (AFM), F5's high-performance, stateful, full-proxy network firewall designed to guard data centers against incoming threats that enter the network can help secure and protect your Exchange deployment.
- The BIG-IP LTM 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.

Products and versions

Product	Version	
Microsoft Exchange Server	2010, 2010 SP1, SP2, and SP3; 2013, 2013 SP1, and all Cumulative Updates (CUs) for those versions	
BIG-IP system	Manual configuration: 11.0 - 11.6 App template: 11.3 - 11.6	
BIG-IP iApp template	f5.microsoft_exchange_2010_2013_cas.v1.5.0	
Deployment Guide version	1.3 See <u>Document Revision History on page 119</u> for revision details	
Last updated	08-06-2015	

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

> For previous versions of this and other guides, see the Deployment guide <u>Archive tab</u> on f5.com: https://f5.com/solutions/deployment-guides/archive-608



Contents

Introduction	3
What is F5 iApp?	3
Prerequisites and configuration notes	4
iApp Deployment Scenarios	7
This BIG-IP LTM will load balance and optimize Client Access Server traffic	7
This BIG-IP LTM will receive HTTP-based Client Access traffic forwarded by a BIG-IP APM	8
This BIG-IP APM will provide secure remote access to CAS	9
Preparation worksheets	10
Configuring the BIG-IP system for Microsoft Exchange using the iApp template	12
Downloading and importing the new iApp	12
Upgrading from a previous version of the iApp template	12
Getting started with the Exchange iApp template	13
Configuring the LTM to receive HTTP-based Client Access traffic forwarded by an APM	34
Configuring the BIG-IP APM to provide secure remote access to Client Access Servers	46
Modifying the iApp configuration	52
Optional: Configuring the BIG-IP system to support MAPI over HTTP in Exchange 2013 SP1	53
Optional: Configuring APM to Support Windows Integrated Authentication For Outlook Web App	55
Optional: Configuring BIG-IP LTM/APM to support NTLMv2-only deployments	56
Troubleshooting	59
Appendix A: Configuring additional BIG-IP settings	65
Appendix B: Using X-Forwarded-For to log the client IP address	67
Appendix C: Manual configuration tables	69
Configuration table if using a combined virtual server for Exchange HTTP-based services	69
Configuration table if using separate virtual servers for Exchange HTTP-based services	71
BIG-IP APM manual configuration	86
Optional: Securing Access to the Exchange 2013 Administration Center with BIG-IP APM	98
Optional: Configuring the APM for Outlook Anywhere with NTLM Authentication - BIG-IP v11.3 or later only	102
Manually configuring the BIG-IP Advanced Firewall Module to secure your Exchange deployment	108
Appendix D: Technical Notes	113
Appendix E: Active Directory and Exchange Server configuration for NTLM	115
BIG-IP APM/LTM without DNS lookups	117
Document Revision History	119

Introduction

This document provides guidance for using the updated, downloadable BIG-IP iApp Template to configure the Client Access server role of Microsoft Exchange Server, as well as instructions on how to configure the BIG-IP system manually. This iApp template was developed for use with both Exchange Server 2013 and 2010.

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, Outlook Anywhere (RPC over HTTP, including the Offline Address Book), ActiveSync, Autodiscover, RPC Client Access (MAPI) for Exchange 2010 only, POP3, IMAP4, and MAPI over HTTP.

For more information on the Client Access Server role, see

- 2010: http://technet.microsoft.com/en-us/library/bb124915%28EXCHG.140%29.aspx
- 2013: http://technet.microsoft.com/en-us/library/bb124558%28v=exchg.150%29.aspx

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

You can also see the BIG-IP deployment guide for SMTP services at: http://www.f5.com/pdf/deployment-guides/f5-smtp-dg.pdf.

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/.

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

What is F5 iApp?

New to BIG-IP version 11, F5 iApp is a powerful set of features in the BIG-IP system that provides a new way to architect application delivery in the data center. iApp includes a holistic, application-centric view of how applications are managed and delivered inside, outside, and beyond the data center. The iApp template for Microsoft Exchange Server acts as the single-point interface for building, managing, and monitoring the Exchange 2010 and 2013 Client Access role.

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

Skip ahead Advanced

If you are already familiar with the Exchange iApp, you can skip directly to the relevant section after reading the prerequisites:

- Configuring the BIG-IP system for Microsoft Exchange using the iApp template on page 12 if using the iApp template, or
- Appendix C: Manual configuration tables on page 69 if configuring the BIG-IP system manually.

Prerequisites and configuration notes

Use this section for important items you need to know about and plan for before you begin this deployment. Not all items will apply in all implementations, but we strongly recommend you read all of these items carefully.

General BIG-IP system prerequisites

- For this deployment guide, the BIG-IP system *must* be running version 11.0 or later. If you are using a previous version of the BIG-IP system, see the Deployment Guide index on F5.com. The configuration in this guide does not apply to previous versions.
- Most of the configuration guidance in this document is performed on F5 devices. We provide a summary of Exchange configuration steps for reference only; for complete information on how to deploy or configure the components of Microsoft Exchange Server, consult the appropriate Microsoft documentation. F5 cannot provide support for Microsoft products.
- > If deploying BIG-IP APM features, you must fully license and provision APM before starting the iApp template.
- This document provides guidance on using the Exchange iApp template. Additionally, for users familiar with the BIG-IP system, there are manual configuration tables at the end of this guide. Because of the complexity of this configuration, we strongly recommend using the iApp to configure the BIG-IP system.
- > F5's advanced health monitors for Autodiscover, Exchange Web Services, and Outlook Anywhere support Basic and NTLMv1 authentication only. See *Troubleshooting on page 59* for more information.

iApp template prerequisites and notes

> This document provides guidance on using the F5 supplied downloadable iApp template for Microsoft Exchange 2010 and 2013 available via downloads.f5.com. The latest official release can always be found at: http://support.f5.com/kb/en-us/solutions/public/13000/400/sol13497.html.

You must use a downloadable iApp for BIG-IP versions 11.0 and later. For the iApp template, you must be using version 11.3 or later as it contains a number of fixes and enhancements not found in the default iApp, or other downloadable versions.



Warning To run the Microsoft Exchange iApp template, you must be logged into the BIG-IP system as a user that is assigned the admin role. For more information on roles on the BIG-IP system, see the BIG-IP User Accounts chapter of the BIG-IP TMOS: Concepts guide.

- If you have an existing Exchange application service from a previous version of the downloadable iApp, see <u>Upgrading</u> from a previous version of the iApp template on page 12 for instructions on how to upgrade the configuration.
- ➤ Exchange 2013 SP1 introduces the MAPI over HTTP transport protocol (http://technet.microsoft.com/en-us/library/dn635177(v=exchg.150).aspx). The iApp template does not yet support this new protocol. See Optional: Configuring the BIG-IP system to support MAPI over HTTP in Exchange 2013 SP1 on page 53 for instructions on configuring the BIG-IP system for MAPI over HTTP.
- If you have existing, manually created Node objects on the BIG-IP system and given these nodes a name, you cannot use the IP addresses for those nodes when configuring the iApp. You must first manually delete those nodes and re-add them without a name, or delete the nodes and let the iApp automatically create them.
- > For some configuration objects, such as profiles, the iApp allows you to import custom objects you created outside the template. This enables greater customization and flexibility. If you have already started the iApp template configuration and then decide to you want to create a custom profile, you can complete the rest of the template as appropriate and then reenter the template at a later time to select the custom object. Otherwise you can exit the iApp immediately, create the profile, and then restart the iApp template from the beginning.

SSL certificate and key prerequisites and notes

▶ If you are using the BIG-IP system to offload SSL (Exchange 2010 and Exchange 2013 SP1 and later only) or for SSL Bridging, we assume you have already obtained an SSL certificate and key, and it is installed on the BIG-IP LTM system. To configure your Client Access servers to support SSL offloading, you must first follow the Microsoft documentation. See http://social.technet.microsoft.com/wiki/contents/articles/how-to-configure-ssl-offloading-in-exchange-2010.aspx.

Make sure you follow the correct steps for the version of Exchange Server that you are using.

- While SSL offload was not supported in the RTM version of Exchange Server 2013, it is now supported in 2013 SP1 (http://social.technet.microsoft.com/wiki/contents/articles/15946.how-to-configure-ssl-offloading-in-exchange-2013.aspx).
 - If you using Exchange 2013 and are not yet on SP1, you must change the default setting for Outlook Anywhere on each Client Access Server so that SSL offloading is not configured.
- ➤ For Exchange Server 2010 and 2013 SP1 only: We generally recommend that you do not re-encrypt traffic between your BIG-IP APM and BIG-IP LTM because both BIG-IP systems must process the SSL transactions. However, if you choose to re-encrypt, we strongly recommend you use a valid certificate (usually SAN-enabled) rather than the default, self-signed certificate for the Client SSL profile on your BIG-IP LTM system. If not re-encrypting traffic, you do not need a certificate on your BIG-IP LTM.
- > This template currently only supports the use of a single DNS name and corresponding certificate and key for all services, or multiple DNS names using a SAN-enabled certificate and key.
- If using a single virtual server for all HTTP-based Client Access services as recommended, you **must** obtain the Subject Alternative Name (SAN) certificate (or wildcard certificate, see the next paragraph) and key from a 3rd party certificate authority that supports SAN certificates, and then import it onto the BIG-IP system. In versions prior to 11.1, the BIG-IP system does not display SAN values in the web-based Configuration utility, but uses these certificates correctly.
 - While the BIG-IP system supports using a wildcard certificate to secure Exchange CAS deployments using multiple FQDNs, for increased security, F5 recommends using SAN certificate(s) where possible. Additionally, some older mobile devices are incompatible with wildcard certificates. Consult your issuing Certificate Authority for compatibility information.
 - ▶ Note: For more information on SAN certificates, see Subject Alternative Name (SAN) SSL Certificates on page 113.

BIG-IP Access Policy Manager prerequisites and notes

- ➤ If you want to display the computer type (public/shared vs private) and light version (Use the light version of Outlook Web App) options for OWA on the APM logon page via the BIG-IP APM, you must run the following PowerShell command on one of your Client Access Servers (only one): Get-OwaVirtualDirectory | Set-OwaVirtualDirectory -LogonPageLightSelectionEnabled \$true -LogonPagePublicPrivateSelectionEnabled \$true
- ➤ If you are using BIG-IP APM, the following table shows the 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	Exchange 2010 and 2013 SP1 only. Optional but strongly recommended
Client Access Array	Not configured	Enabled	Exchange 2010 only: Required
OWA Authentication ¹	Forms ²	Forms (default) ² or Windows authentication (smart card)	Required
Autodiscover Authentication ¹	Negotiate	Negotiate (default)	Required
ActiveSync Authentication¹	Basic	Basic (default)	Required
Outlook Anywhere Authentication ^{1,3}	2010:Basic 2013: Negotiate	Basic (default) or NTLM	Required

¹ Exchange Server 2010 and 2013 SP1 and later only. 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 later in this guide.

³ Outlook Anywhere is disabled by default in Exchange 2010; you must enable it before you can use it. You can optionally configure BIG-IP APM v11.3 and later for NTLM authentication for Outlook Anywhere. See page 50.

(i) Important The values in the following table are only examples, use the values appropriate for your configuration.

In our example, we use the following conventions.

Role	FQDNs	DNS Records	External URL/ Host name	Notes
	Combined virtual server			If the external DNS SRV
	mail.example.com	A: mail.example.com SRV: _autodiscovertcp.example.com: port 443, Host 'mail.example.com.' https://mail.example.com		record listed is not used, and you don't want to use SCP internally, you must also have
Autodiscover	Separate virtual servers			at least one of these, set to the same IP as your OWA FQDN: example.com autodiscover.example.com
	autodiscover.example.com	A: autodiscover.example.com SRV: _autodiscovertcp.example.com: port 443, Host 'autodiscover.example.com.'	https://autodiscover.example.com/ autodiscover/autodiscover.xml	
	Combined virtual server			
0 11 1 11 1 1	mail.example.com	A: mail.example.com	https://mail.example.com/owa	
Outlook Web App	Separate virtual servers			
	owa.example.com	A: owa.example.com	https://owa.example.com/owa	
	Combined virtual server			
	mail.example.com	A: mail.example.com	https://mail.example.com/ Microsoft-Server-ActiveSync	
ActiveSync	Separate virtual servers			
	mobile.example.com	A: mobile.example.com	https://mobile.example.com/ Microsoft-Server-ActiveSync	
	Combined virtual server			To prevent internal users from
	mail.example.com	A: mail.example.com	mail.example.com	receiving a password prompt, your internal DNS must not
	Separate virtual servers	have an A record for the FQDN for Outlook Anywhere. This		
Outlook Anywhere (RPC over HTTP)	oa.example.com	A: oa.example.com	oa.example.com	only applies if you are using Exchange 2010, using RPC MAPI internally and Outlook Anywhere externally, and your internal clients do not have a route to the external Outlook Anywhere/EWS virtual server(s).
	Combined virtual server			
Outlook Anywhere	mail.example.com	A: mail.example.com	https://mail.example.com/mapi	
(MAPI over HTTP)	Separate virtual servers			
	mapi.example.com	A: mapi.example.com	https://mapi.example.com/mapi	
	Separate virtual servers			
RPC Mapi ¹	array.example.com	A: array.example.com	N/A	
	*		·	·

¹ Exchange Server 2010 only. Exchange 2013 does not use RPC.

For more information, see:

- 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

iApp Deployment Scenarios

The iApp greatly simplifies configuring the BIG-IP system for Microsoft Exchange 2010/2013 Client Access Server roles. Before beginning the Application template, you must make a decision about the scenario in which you are using BIG-IP system for this deployment. The iApp presents the following three deployment options. You choose one of these options when you begin configuring the iApp.

- This BIG-IP LTM will load balance and optimize Client Access Server traffic, on this page
- This BIG-IP LTM will receive HTTP-based Client Access traffic forwarded by a BIG-IP APM on page 8
- This BIG-IP APM will provide secure remote access to CAS on page 9

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

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

You would not select this option if you intend to deploy a separate APM that provides secure remote access to CAS HTTP services.

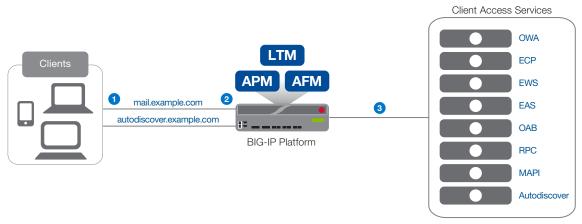


Figure 1: Logical configuration example showing the BIG-IP system directing traffic to Client Access Services

The traffic flow for this scenario is:

- 1. All Exchange Client Access traffic goes to the BIG-IP system.
- 2. You can use the following optional modules if they are licenced and provisioned on you BIG-IP system:
 - BIG-IP Access Policy Manager (APM)
 The BIG-IP APM module provides secure access and proxied authentication (pre-authentication) for HTTP-based Client Access services: Outlook Web App, Outlook Anywhere, ActiveSync, and Autodiscover). The BIG-IP APM presents a login page to end users that takes the place of the forms-based login page normally presented by Outlook Web App. Users provide credentials through the BIG-IP APM form; the BIG-IP APM then authenticates the user to Active Directory.
 - BIG-IP Advanced Firewall Manager (AFM)
 The BIG-IP Advanced Firewall Manager (AFM) is a high-performance, stateful, full-proxy network firewall designed to guard data centers against incoming threats that enter the network on the most widely deployed protocols.
- 3. The BIG-IP LTM load balances and optimizes the traffic to the Client Access Servers, including the services which are not HTTP-based: RPC Client Access, POP3, and IMAP4.

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

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

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

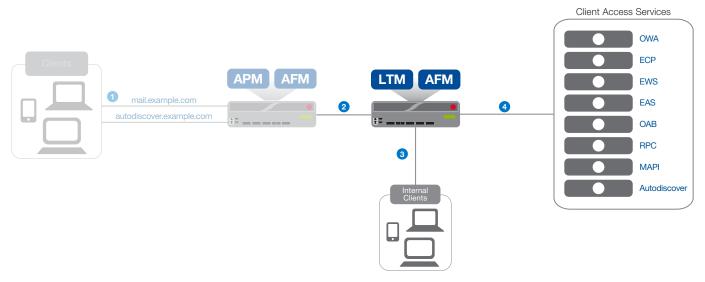


Figure 2: Logical configuration example showing the BIG-IP system receiving traffic from a BIG-IP APM

- 1. Traffic comes in from the BIG-IP APM as described in the next scenario.
- 2. The BIG-IP LTM receives HTTP-based Client Access traffic from a separate BIG-IP APM, or directly received the non HTTP-based traffic.
- 3. If you have internal Exchange clients, all Client Access Server traffic from the internal clients goes directly to the BIG-IP LTM.
- 4. The BIG-IP LTM load balances and optimizes the traffic to the Client Access Servers, including the services which are not HTTP-based: RPC Client Access (MAPI), POP3, and IMAP4.
- Note: While this scenario can accommodate internal clients, we do not recommend using this virtual server in that way. We strongly recommend creating a second instance of the iApp on this BIG-IP LTM for the direct traffic/internal users. You must use a unique virtual server IP address; all of the other settings can be identical. Once both iApps have been created, you would configure Split DNS (use the same domain name, but different zones and IP addresses for internal and external clients). For more information about Split DNS, refer to your DNS documentation.

This BIG-IP APM will provide secure remote access to CAS

You can select this scenario to configure the BIG-IP system as a BIG-IP APM that will use a single virtual server to provide proxy authentication (pre-authentication) and secure remote access to Exchange HTTP-based Client Access services without requiring the use of an F5 Edge Client. When you select this deployment scenario, the BIG-IP APM presents a login page to end users that takes the place of the forms-based login page normally presented by Outlook Web App. Users provide credentials through the BIG-IP APM form: the

BIG-IP APM then authenticates the user to Active Directory. The BIG-IP system will only forward connections after a user has authenticated successfully. The traffic is then sent to another BIG-IP running LTM which provides advanced load balancing, persistence, monitoring and optimizations for HTTP-based Client Access services.

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

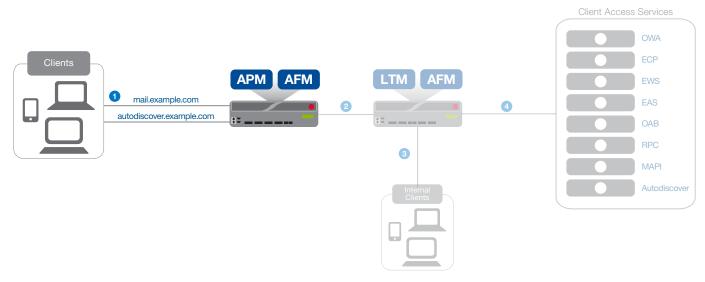


Figure 3: Logical configuration example showing the BIG-IP APM providing proxy authentication and secure remote access

- 1. HTTP-based Client Access traffic goes to the BIG-IP APM, which provides proxy authentication and secure remote access.
 - Note: While this scenario can accommodate internal clients, we do not recommend using this virtual server in that way. We strongly recommend creating a second instance of the iApp on this BIG-IP LTM for the direct traffic/internal users. You must use a unique virtual server IP address; all of the other settings can be identical. Once both iApps have been created, you would configure Split DNS (use the same domain name, but different zones and IP addresses for internal and external clients). For more information about Split DNS, refer to your DNS documentation.
- 2. After authentication, the BIG-IP APM sends the traffic to a separate BIG-IP LTM for intelligent traffic management.

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

Preparation worksheets

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

	BIG-IP LTM Preparat	ion workshe	eet
Traffic arriving to this BIG-IP	Encrypted		Unencrypted
system is:	SSL Certificate: Key: If re-encrypting traffic to the Client Access Server using the BIG-IP default certificate and key for the profile: Certificate:	Server SSL	If encrypting traffic to the Client Access Servers and not using the BIG-IP default certificate and key: Certificate: Key:
	Key:		
BIG-IP virtual servers and	Same Subnet		Different Subnets
Client Access Servers will be on:	If the maximum number of expected concurrent under the Client Access Server is more than 6,000, you neel IP address for each 6,000 users or fraction thereof the concurrent under	d one SNAT of:	If the CAS servers are a different subnet from the BIG-IP virtual servers, and do not use the BIG-IP as their default gateway, and if the maximum number of expected concurrent users per Client Access Server is more than 6,000, you need one SNAT IP address for each 6,000 users or fraction thereof: 1:
Single virtual IP address for	Single virtual IP address		Different virtual IP addresses for different services
all Client Access Services or multiple addresses	IP address for the BIG-IP virtual server:		You need a unique IP address for each of the Client Access services you are deploying: Outlook Web App: Outlook Anywhere: ActiveSync: Autodiscover: RPC Client Access (MAPI) - 2010 only: POP3: IMAP4:
Are all Client Access services	Same set of Client Access Servers for all	services	Different Client Access Servers for different services
handled by the same set of servers, or different Servers for different services?	IP addresses of the Client Access Servers: 1: 6:		IP addresses for Client Access Servers for each service: Outlook Web App:
	2: 7:		Outlook Anywhere:
	3: 8:		ActiveSync:
	4: 9:		Autodiscover:
	5: 10:		RPC Client Access (2010):
			POP3:
			IMAP4:
RPC Client Access ports (Exchange 2010 only)	If you are deploying RPC Client Access (MAPI), ar MAPI:		t to use the default Dynamic port range, specify a port for:
	Address Book:		

	BIG-IP LTM Preparation workshe	et	
Advanced Monitor configuration	If you want the iApp to configure advanced health monitors which perform logins to HTTP-based, POP3, and IMAP4 Client Access services (as opposed to simple monitors which only check network connectivity), you need the following information:		
	If deploying Autodiscover, email address for monitoring:	Second mailbox for monitoring (recommended): If deploying Autodiscover, 2 nd email address for monitoring:	
	Mailbox account name in Active Directory for the monitors:	2 nd mailbox account name in Active Directory for the monitors:	
	Associated password:	Associated password for this account:	
	Domain name (can be FQDN or NETBIOS) of the user account used for monitors:	2 nd domain name (can be FQDN or NetBIOS) of the user account used for monitors:	
	Important: Advanced monitors for Autodiscover, EWS, and Outlook Anywhere support Basic and NTLMv1 authentication only.		
Outlook Web App authentication method	If deploying Outlook Web App, which authentication method have you configured: Forms-Based Authentication (default) Important: If you are deploying BIG-IP APM, you must use Forms-Based. Basic or Windows Integrated authentication		
Same FQDN for all HTTP-	Same FQDN	Different FQDNs	
based Client Access services or different FQDNs	FQDN for all HTTP-based Client Access services:	You need a FQDN for each HTTP-based Client Access services you are deploying:	
		Outlook Web App:	
		Outlook Anywhere:	
		ActiveSync:	
		Autodiscover:	

BIG-IP Access Policy Manager Preparation Worksheet		
Outlook Web App FQDN	If you are deploying APM and OWA, you need the FQDN this is used to access OWA (such as owa.example.com):	
Domain Controller FQDNs and IP addresses that the BIG-iP system can contact	What are the Domain Controller FQDNs and IP address this BIG-IP system can contact (use FQDN and not NETBIOS name) 1:	
Active Directory Domain name for Exchange users	What is the Active Directory Domain name (must be in FQDN format):	
Active Directory Anonymous binding	If Anonymous Binding is not allowed in your Active Directory implementation, you need an Active Directory account with administrative permissions: User name: Password:	
If deploying the "BIG-IP APM w	rill provide secure remote access to CAS" scenario	
BIG-IP APM virtual server	What is the IP address you want to use for your BIG-IP APM virtual server:	
SSL Certificate and Key	SSL Certificate:	
Re-encrypt the traffic to the BIG-IP virtual server	You must know if the remote BIG-IP LTM that will receive traffic from this BIG-IP APM is using a self-signed/default certificate and key or a certificate signed by a Certificate Authority.	
Remote LTM virtual server	What is the virtual server address on the remote BIG-IP LTM to which this BIG-IP APM will forward traffic:	

	BIG-IP Advanced Firewall Manager Preparation Worksheet
Subnets/Networks	Which networks or subnets should be allowed to access the Exchange deployment:

Configuring the BIG-IP system for Microsoft Exchange using the iApp template

Use this section for guidance on configuring the BIG-IP system using the iApp template. If you plan to configure the system manually, see Appendix C: Manual configuration tables on page 69.

Downloading and importing the new iApp

The first task is to download and import the new Exchange Server Client Access Server iApp template.

To download and import the iApp

- 1. Open a web browser and go to http://support.f5.com/kb/en-us/solutions/public/13000/400/sol13497.html.
- Follow the instructions to download the Microsoft Exchange iApp to a location accessible from your BIG-IP system.
- Extract (unzip) the f5.microsoft_exchange_2010_2013_cas<latest version>.tmpl file. 3.
- 4 Log on to the BIG-IP system web-based Configuration utility.
- 5. On the Main tab, expand iApp, and then click Templates.
- Click the **Import** button on the right side of the screen. 6.
- Click a check in the **Overwrite Existing Templates** box.
- 8. Click the **Browse** button, and then browse to the location you saved the iApp file.
- Click the **Upload** button. The iApp is now available for use.

Upgrading from a previous version of the iApp template

If you configured your BIG-IP system using a previous version of the downloadable iApp template, we strongly recommend you upgrade the iApp template to this current version. You cannot upgrade an application service that was based on the f5.microsoft. exchange 2010 template; you can only upgrade if you used one of the downloadable iApp templates.

When you upgrade the template, you simply change the parent template on the application service you previously created, and then make any necessary modifications to take advantage of new functionality.

(i) Important If you are upgrading from an iApp version prior to v.1.2.0, carefully review all settings before submitting the template. For example, if you had configured the original template for SSL bridging, after upgrading this setting defaults back to SSL offload, and you must change it.

To upgrade an existing application service to the new iApp template

- From the Main tab of the BIG-IP Configuration utility, expand iApp and then click Application Services.
- 2. Click the name of your existing Microsoft Exchange application service from the list.
- 3. On the Menu bar, click Reconfigure.
- At the top of the page, in the Template row, click the Change button to the right of the list.
- From the Template list, select f5.microsoft_exchange_2010_2013_cas.<latest version>. 5.
- Review the questions in the new template, making any necessary modifications. Use the iApp walkthrough section of this guide for information on specific questions.

If you used a previous version of the iApp to deploy BIG-IP APM, you must add the FQDN and IP address for each Active Directory server in your domain that the BIG-IP system can contact. The iApp now creates a pool for the Active Directory servers (even for only one server), where in previous versions of the template you could only specify a single Active Directory server.

Click Finished.



Marning The option for restricting EAC access by IP address or network is no longer available in iApp v1.4.0 (beginning in v1.4.0rc2) as it did not function reliably. If you are upgrading from v1.3.0 of the iApp template, and enabled that option, it will no longer be a part of the configuration after the upgrade. The option to have APM restrict EAC access to members of the Exchange Organization Management Security Group is still available.

Getting started with the Exchange iApp template

To begin the Exchange iApp Template, use the following procedure.

To start the iApp template

- 1. Log on to the BIG-IP system.
- 2. On the Main tab, expand iApp, and then click Application Services.
- 3. Click Create. The Template Selection page opens.
- 4. In the Name box, type a name. In our example, we use Exchange-2013_.
- From the Template list, select f5.microsoft_exchange_2010_2013_cas.<latest version>.
 The new Microsoft Exchange template opens.

Advanced options

If you select **Advanced** from the **Template Selection** list at the very top of the template, you see Device and Traffic Group options for the application. This feature, introduced in v11.0, is a part of the Device Management configuration. This functionality extends the existing High Availability infrastructure and allows for clustering, granular control of configuration synchronization and granular control of failover. To use the Device and Traffic Group features, you must have already configured Device and Traffic Groups before running the iApp. For more information on Device Management, see the product documentation.

1. Device Group

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

2. Traffic Group

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

Inline help

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

Do you want to see inline help?

Select whether you want to see informational and help messages inline throughout the template, or if you would rather hide this inline help. If you are unsure, we recommend having the iApp display all inline help.

Important and critical notes are always shown, no matter which selection you make.

• Yes, show inline help text

Select this option to see all available inline help text.

· No, do not show inline help

If you are familiar with this template, or with the BIG-IP system in general, select this option to hide the inline help text.

Deployment Scenario

Choose the option that best describes how you plan to use the BIG-IP system you are currently configuring. The scenario you select from the list determines the questions that appear in the rest of the iApp. The scenarios were described in <u>iApp Deployment Scenarios on page 7</u>.

1. Which scenario describes how you will use the BIG-IP system?

Choose the scenario that best describes the way you plan to use this BIG-IP system. Guidance for each scenario is contained in a separate section of this deployment guide. Click the link to go to the relevant section of the guide for the scenario you plan to deploy.

BIG-IP LTM will load balance and optimize CAS traffic

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

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

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

For this role, go to Configuring the BIG-IP LTM to load balance and optimize Client Access Server traffic on page 15.

BIG-IP LTM will receive HTTP-based CAS traffic forwarded by a BIG-IP APM

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

For this role, go to Configuring the LTM to receive HTTP-based Client Access traffic forwarded by an APM on page 34.

BIG-IP APM will provide secure remote access to CAS

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

For this role, go to Configuring the BIG-IP APM to provide secure remote access to Client Access Servers on page 46.

2. Which version of Exchange are you using?

Choose the version of Microsoft Exchange Server you are using. Some features of the iApp are available only to a particular version.

• Exchange Server 2010

Select this option if you are deploying the BIG-IP system for Microsoft Exchange 2010.

Exchange Server 2013

Select this option if you are deploying the BIG-IP system for Microsoft Exchange 2013.

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

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

Analytics

This section of the template asks questions about Analytics. The Application Visibility Reporting (AVR) module allows you to view statistics specific to your Microsoft Exchange implementation. AVR is available on all BIG-IP systems v11 and later, however you must have the AVR provisioned for this option to appear. Note that this is only for application visibility reporting, you can view object-level statistics from the BIG-IP without provisioning AVR.

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

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

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

Select whether you want to enable AVR for Analytics for HTTP-based services. Note that Analytics does not always properly report the HTTP methods of Outlook Anywhere.

• No, do not enable Analytics

Select this option if you do not want to use Analytics, and then continue with BIG-IP Access Policy Manager.

• Yes, enable Analytics using AVR

If you choose to enable Analytics, select Yes from the list, and then answer the following questions.

a. <u>Use the default Analytics profile or select a custom profile?</u>

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

• Select a custom Analytics profile

Select this option if you have already created a custom Analytics profile for Exchange Server.

a. Which Analytics profile do you want to use?

From the list, select the appropriate Analytics profile.

Use default profile

Select this option if you have not yet created a custom Analytics profile for Microsoft Exchange. We do not recommend using the default profile.

BIG-IP Access Policy Manager

This section in this scenario asks about BIG-IP APM. To use APM, it must be fully licensed and provisioned before starting the template. If you are not deploying BIG-IP APM, continue with the next section. As mentioned in the prerequisites, if you are deploying APM, you must have configured the BIG-IP system for DNS and NTP; see <u>Configuring DNS and NTP settings on page 65</u> for instructions.

1. Provide secure authentication to CAS HTTP-based services with BIG-IP Access Policy Manager?

Specify whether you want to deploy BIG-IP APM to provide proxy authentication and secure remote access for HTTP-based Client Access services.

No, do not provide secure authentication using BIG-IP APM

Select this option if you do not want to use the BIG-IP APM at this time. You can always reconfigure the iApp template at a later date should you decide to add BIG-IP APM functionality.

Yes, provide secure authentication using BIG-IP APM

Select this option if you want to use the BIG-IP APM to provide proxy authentication and secure remote access for your Exchange deployment.

a. Would you like to create a new Access Profile, or use an existing Access Profile?

Choose whether you want the system to create a new BIG-IP APM Access Profile, or if you have already created a custom Access Profile outside the template. If you are unsure, select **Create a new Access Profile**.

Select the Access profile you created from the list

If you have previously created an Access profile for your Exchange implementation, select the existing profile you created from the list. Continue with the next section.

Create a new Access profile

Select this option if you have not created a custom Access profile, and want the system to create one.

a. Would you like to create a new AAA server, or use an existing AAA server?

Choose whether you want the system to create a new BIG-IP APM AAA Server object, or if you have already created a custom AAA Server outside the template. The AAA server contains information about your Active Directory implementation. If you are unsure, select **Create a new AAA Server**.

Select the AAA Server you created from the list

If you have previously created an AAA Server for your Exchange implementation, select the existing object you created from the list.

a. What is the FQDN of your Active Directory domain for your Exchange users?

Specify the FQDN of the Active Directory deployment for your Exchange users. This is the FQDN for your entire domain, such as example.com, rather than the FQDN for any specific host. Continue with the *What text should appear in the user access logon prompt* question on the following page.

Create a new AAA Server

Select this option if you have not created a custom AAA Server, and want the system to create one.

a. What is the FQDN of your Active Directory domain for your Exchange users?

Specify the FQDN of the Active Directory deployment for your Exchange users. This is the FQDN for your entire domain, such as example.com, rather than the FQDN for any specific host.

b. Which Active Directory servers in your domain can this BIG-IP system contact?

Specify both the FQDN and IP address of each Active Directory server you want the BIG-IP APM to use for servicing authentication requests. Click **Add** to include additional servers.

c. <u>Does your Active Directory domain allow anonymous binding?</u>

Select whether anonymous binding is allowed in your Active Directory environment.

· Yes, anonymous binding is allowed

Select this option if anonymous binding is allowed. No further information is required.

· No, credentials are required for binding

If credentials are required for binding, you must specify an Active Directory user name and password for use in the AAA Server.

a. Which Active Directory user with administrative permissions do you want to use? Type a user name with administrative permissions.

b. What is the password associated with that account?

Type the associated password.

d. How do you want to handle health monitoring for this pool?

Specify whether you want the template to create a new LDAP monitor or a new ICMP monitor, or if you select an existing monitor. For more accurate monitoring, we recommend using an LDAP monitor.

Select an existing monitor for the Active Directory pool

Select this option if you have already created a health monitor (only monitors with a Type of LDAP or External can be used) for the Active Directory pool that will be created by the template. If you want to create a health monitor, but have not already done so, you must exit the template and create the object before it becomes available from the list.

The iApp allows you to select monitors that are a part of another iApp Application Service. If you select a monitor that is a part of another Application Service, be aware that any changes you make to the monitor in the other Application Service will apply to this Application Service as well.

a. Which monitor do you want to use?

From the list, select the LDAP or External monitor you created to perform health checks for the Active Directory pool created by the template. Only monitors that have a Type value of LDAP or External appear in this list. Continue with the "What text should appear in the user access logon prompt" question on this page.

• Use a simple ICMP monitor for the Active Directory pool

Select this option if you only want a simple ICMP monitor for the Active Directory pool. This monitor sends a ping to the servers and marks the server UP if the ping is successful. Continue with the "What text should appear in the user access logon prompt" question on this page.

Create a new LDAP monitor for the Active Directory pool

Select this option if you want the template to create a new LDAP monitor for the Active Directory pool. You must answer the following questions:

a. Which Active Directory user name should the monitor use?

Specify an Active Directory user name for the monitor to use when attempting to log on as a part of the health check. This should be a user account created specifically for this health monitor, and *must* be set to never expire.

b. What is the associated password?

Specify the password associated with the Active Directory user name.

c. What is the LDAP tree for this user account?

Specify the LDAP tree for the user account. As noted in the inline help, ADSI editor, an tool for Active Directory LDAP administration, is useful for determining the correct LDAP tree value. For example, if the user name is 'user1' which is in the organizational unit 'Exchange Users' and is in the domain 'exchange.example.com', the LDAP tree would be: ou=Exchange Users, dc=Exchange, dc=example, dc=com.

d. Does your Active Directory domain require a secure protocol for communication?

Specify whether your Active Directory implementation requires SSL or TLS for communication, or does not require a secure protocol. This determines the port the health monitor uses.

No, a secure protocol is not required

Select this option if your Active Directory domain does not require a secure protocol.

· Yes, SSL communication is required

Select this option if your Active Directory domain requires SSL communication. The health check uses port 636 as the Alias Service Port.

· Yes, TLS communication is required

Select this option if your Active Directory domain requires TLS communication. The health check uses port 389 as the Alias Service Port.

e. How many seconds between Active Directory health checks?

Specify how many seconds the system should use as the health check Interval for the Active Directory servers. We recommend the default of 10 seconds.

b. What text should appear in the user access logon prompt?

Type the text you want users to see above the user name and password prompts when logging on to the BIG-IP APM. By default, this includes the HTML
br> tag to insert a line break between 'Secure Logon' and 'for F5 Networks'.

Application Firewall Manager (BIG-IP AFM)

This section gathers information about BIG-IP Advanced Firewall Manager if you want to use it to protect the Exchange deployment. For more information on configuring BIG-IP AFM, see http://support.f5.com/kb/en-us/products/big-ip-afm.html, and then select your version.

1. Do you want to use BIG-IP AFM to protect your application?

Choose whether you want to use BIG-IP AFM, F5's network firewall, to secure this Exchange deployment. If you choose to use BIG-IP AFM, you can restrict access to the Exchange virtual server(s) to a specific network or IP address. See the BIG-IP AFM documentation for specific details on configuring AFM.

No, do not use Application Firewall Manager

Select this option if you do not want to enable BIG-IP AFM at this time. You can always re-enter the template at a later date to enable BIG-IP AFM. Continue with the next section.

Select an existing AFM policy from the list

If you already created a BIG-IP AFM policy for this implementation, select it from the list. Continue with c.

▶ Yes, use F5's recommended AFM configuration

Select this option if you want to enable BIG-IP AFM using F5's recommended configuration.

a. Do you want to restrict access to your application by network or IP address?

Choose whether you want to restrict access to the Exchange implementation via the BIG-IP virtual server.

▶ No, do not restrict source addresses (allow all sources)

By default, the iApp configures the AFM to accept traffic destined for the Exchange virtual server(s) from all sources. If you do not have a need to restrict access to the virtual server, leave this option selected and then continue with \boldsymbol{b} .

Restrict source addresses

Select this option if you want to restrict access to the Exchange virtual server(s) by IP address or network address.

i). What IP or network addresses should be allowed to access your application?

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

b. How do you want to control access to your application from sources with a low reputation score?

The BIG-IP AFM uses an IP intelligence database to categorize IP addresses coming into the system. Choose what you want the system to do for sources that are attempting to access the Exchange virtual server(s) with a low reputation score. For more information, see the BIG-IP AFM documentation.

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

Allow all sources regardless of reputation

Select this option to allow all sources, without taking into consideration the reputation score.

▶ Reject access from sources with a low reputation

Select this option to reject access to the Exchange virtual server(s) from any source with a low reputation score.

▶ Allow but log access from sources with a low reputation

Select this option to allow access to the Exchange virtual server(s) from sources with a low reputation score, but add an entry for it in the logs.

c. Would you like to stage a policy for testing purposes?

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

Do not apply a staging policy

Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.

Select an existing policy from the list

If you have already created a firewall policy for this implementation, select it from the list. Only policies that already exist on the system appear in the list. To create a new policy, on the Main tab, click **Security > Network Firewall > Policies**. Specific instructions for creating a firewall policy is outside the scope of this iApp and deployment guide.

d. Which logging profile would you like to use?

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

Do not use a logging profile

Select this option if you do not want to use a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.

▶ Select an existing logging profile from the list

If you have already created a logging profile for this implementation, select it from the list. You must create a profile before it is available in the list. To create a logging profile, on the Main tab, click Security > Event Logs > Logging Profiles. Specific instructions for creating a logging profile is outside the scope of this iApp and deployment guide. See the online help or the About Local Logging with the Network Firewall chapter of the BIG-IP Network Firewall: Policies and Implementations guide for more information.

Tell us about your deployment

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

1. Will incoming traffic arrive at this BIG-IP system encrypted or unencrypted?

This question does not appear if you chose to deploy APM in the previous section. If you selected to deploy APM, continue with the re-encrypt question (a) under Encrypted.

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

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

Encrypted

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

- a. Do you want to re-encrypt this traffic to your Client Access Servers?
 - If you are using Exchange 2010 or Exchange 2013 SP1 and later, and want the BIG-IP system to offload SSL processing from the Client Access Servers, select **Do not re-encrypt (SSL Offload)** from the list. Offloading SSL on the BIG-IP system can extend Exchange Server server capacity.
 - (i) Important If you are deploying Exchange Server 2013 and have <u>not</u> installed SP1 or later, you must choose Re-encrypt (SSL Bridging).
 - Do not re-encrypt (SSL Offload)

Select this option if you want to offload SSL processing onto the BIG-IP system. If you choose SSL Offload, you must have followed the instructions described in the prerequisites for configuring the Exchange Server: http://social.technet.microsoft.com/wiki/contents/articles/how-to-configure-ssl-offloading-in-exchange-2010.aspx.

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

The iApp can create a new Client SSL profile, or if you have created a Client SSL profile which contains the appropriate SSL certificate and key for your Exchange implementation, you can select it from the list.

- Select the Client SSL profile you created from the list
 If you manually created a Client SSL profile, select it from the list, and then continue with #2.
- Create a new Client SSL profile
 Select this option if you want the iApp to create a new Client SSL profile.
 - a. Which SSL certificate do you want to use?
 Select the SSL certificate you imported onto the BIG-IP system for decrypting client connections.

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

Note: Any certificate that you obtain with multiple names must be in SAN (Subject Alternative Name) format, not SNI (Server Name Indication) format. For more information on SAN certificates, see Subject Alternative Name (SAN) SSL Certificates on page 113.

b. Which SSL key do you want to use? Select the associated key from the list.

Re-encrypt (SSL Bridging)

Select this option if your implementation requires encrypted traffic to the Client Access Servers, or you are using Exchange 2013. The BIG-IP system unencrypts, then re-encrypts the traffic headed for the Client Access Servers.

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

The iApp can create a new Client SSL profile, or if you have created a Client SSL profile which contains the appropriate SSL certificate and key for your Exchange implementation, you can select it from the list.

• Select the Client SSL profile you created from the list

If you manually created a Client SSL profile, select it from the list, and then continue with #2.

Create a new Client SSL profile

Select this option if you want the iApp to create a new Client SSL profile.

a. Which SSL certificate do you want to use?

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

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

Note: Any certificate that you obtain with multiple names must be in SAN (Subject Alternative Name) format, not SNI (Server Name Indication) format. For more information on SAN certificates, see Subject Alternative Name (SAN) SSL Certificates on page 113.

b. Which SSL key do you want to use?

Select the associated key from the list.

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

Select whether you want the iApp to create an F5 recommended Server SSL profile, or if you want to choose a Server SSL profile you already created.

Important If you are configuring SSL Bridging and using BIG-IP version 11.4.x, you must see When using SSL Bridging and BIG-IP version 11.4.x, pool members may be marked down or you may experience connection resets and TLS errors logged to the Client Access servers on page 62.

Select the Server SSL profile you created from the list

If you have previously created a Server SSL profile for your Exchange implementation, from the list, select the existing Server SSL profile you created.

Create a new Server SSL profile

Select this option if you want the iApp to create a new Server SSL profile.

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

Unencrypted

Select this option if Client Access traffic is arriving at this BIG-IP system unencrypted (typically because you configured to offload SSL/TLS traffic at the BIG-IP APM that is sending Client Access traffic to this device).

a. Do you want to encrypt the traffic to your Client Access Servers?

If you are using Exchange 2010 or Exchange 2013 SP1 and later, and want the BIG-IP system to offload SSL processing from the Client Access Servers, select Do not encrypt (SSL Offload) from the list. Offloading SSL on the BIG-IP system can extend Exchange Server server capacity.

(i) Important If you are deploying Exchange Server 2013 and have not installed SP1 or later, you must choose Encrypt (SSL Bridging).

· Do not encrypt (SSL Offload)

Select this option if you do not want the BIG-IP system to encrypt the traffic destined for the Client Access servers. The BIG-IP system does not modify the traffic, and you can continue with the next question.

• Encrypt (SSL Bridging)

Select this option if your implementation requires encrypted traffic to the Client Access Servers, or you are using Exchange 2013 prior to SP1. The BIG-IP system encrypts the traffic headed for the Client Access Servers.

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

Select whether you want the iApp to create an F5 recommended Server SSL profile, or if you want to choose a Server SSL profile you already created.

i Important If you are configuring SSL Bridging and using BIG-IP version 11.4.x, you must see When using SSL Bridging and BIG-IP version 11.4.x, pool members may be marked down or you may experience connection resets and TLS errors logged to the Client Access servers on page 62.

· Select the Server SSL profile you created from the list

If you have previously created a Server SSL profile for your Exchange implementation, select the existing Server SSL profile you created from the list.

· Create a new Server SSL profile

Select this option if you want the iApp to create a new Server SSL profile.

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

2. How should the system optimize client-side TCP connections to the BIG-IP LTM?

Select how the system should optimize client-side TCP connections. The iApp uses your selection to configure the proper TCP optimization settings on the TCP profile.

• Optimize TCP connections for WAN clients

Select this option if most Exchange server clients are coming into your Exchange environment over a Wide Area Network.

Optimize TCP connections for WAN clients

Select this option if most Exchange server clients are coming into your Exchange environment over a Local Area Network.

3. Where will your BIG-IP virtual servers be in relation to your Client Access Servers?

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

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

Select this option if the BIG-IP virtual servers and the Client Access Servers are on the same subnet. In this case SNAT is configured on the BIG-IP virtual server and you must specify the number of concurrent connections.

a. What is the maximum number of concurrent users you expect per Client Access Server?

Select whether you expect more or fewer than 6,000 concurrent users to each Client Access Server. This answer is used to determine what type of SNAT (secure network address translation) that system uses. A SNAT is an object that maps the source client IP address in a request to a translation address defined on the BIG-IP device.

Note: For specific information on SNAT Pools, including why we chose 6,000 concurrent users per Client Access Server, see Maximum number of concurrent users: SNAT Pool quidance on page 113.

Fewer than 6000

Select this option if you expect fewer than 6,000 concurrent users per Client Access Server. With this option, the system applies SNAT Auto Map, which doesn't require any additional IP addresses, as the system uses an existing self IP address for translation.

More than 6000

Select this option if you expect more than 6,000 users at one time to each server. With this option, the iApp creates a SNAT Pool (or you can choose one you created), for which you need one IP address for each 6,000 users you expect.

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

Select whether you want the system to create a new SNAT Pool, or if you have already created a SNAT pool for this implementation.

• Select the SNAT pool you created from the list

If you have previously created a SNAT Pool for your Exchange implementation, select it from the list.

Create a new SNAT pool

If you have not created a custom SNAT pool, select this option for the iApp to create a new one.

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

Specify one otherwise unused IP address for every 6,000 concurrent connections, or fraction thereof. Click **Add** for additional rows.

(i) Important If you choose more than 6,000 users, but do not specify enough SNAT pool addresses, after the maximum connection limit of 6,000 concurrent users per server is reached, new requests fail.

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

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

a. How have you configured routing on your Client Access Servers?

Select whether the Client Access Servers use this BIG-IP system's Self IP address as their default gateway.

Client Access Servers do NOT use BIG-IP as their default gateway

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

a. What is the maximum number of concurrent users you expect per Client Access Server?

Select whether you expect more or fewer than 6,000 concurrent users to each Client Access Server. This answer is used to determine what type of SNAT that system uses. A SNAT is an object that maps the source client IP address in a request to a translation address defined on the BIG-IP device.

Note: For specific information on SNAT Pools, including why we chose 6,000 concurrent users per Client Access Server, see Maximum number of concurrent users: SNAT Pool guidance on page 113.

Fewer than 6000

Select this option if you expect fewer than 6,000 concurrent users per Client Access Server. With this option, the system applies SNAT Auto Map, which doesn't require any additional IP addresses, as the system uses an existing self IP address for translation.

More than 6000

Select this option if you expect more than 6,000 users at one time to each server. With this option, the iApp creates a SNAT Pool (or you can choose one you created), for which you need one IP address for each 6,000 users you expect.

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

Select whether you want the system to create a new SNAT Pool, or if you have already created a SNAT pool for this implementation.

· Select the SNAT pool you created from the list

If you have previously created a SNAT Pool for your Exchange implementation, select it from the list.

Create a new SNAT pool

If you have not created a custom SNAT pool, select this option for the iApp to create a new one.

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

Specify one otherwise unused IP address for every 6,000 concurrent connections, or fraction thereof. Click Add for additional rows.

(i) Important If you choose more than 6,000 users, but do not specify enough SNAT pool addresses, after the maximum connection limit of 6,000 concurrent users per server is reached, new requests fail.

Client Access Servers use the BIG-IP as their default gateway

Select this option if the Client Access Servers use the BIG-IP system as their default gateway. In this scenario, no additional configuration is necessary to ensure correct server response handling.

4. Will you use a single IP address for all CAS connections, or will you have separate IP addresses?

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

Single IP address

With a single 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 or a wildcard certificate. For detailed information on SAN certificates, see <u>Subject Alternative Name (SAN) SSL Certificates on page 113</u>.

Different IP addresses for different services

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

5. How are you distributing the CAS protocols between servers?

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

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

- All services will be handled by the same set of Client Access Servers

 Choose this option if you are using the same Client Access Servers for all of your Exchange Client Access services.
- Each service will be handled by a unique set of Client Access Servers

 Choose this option if you are using different sets of Client Access Servers for each Client Access service.

Tell us about which services you are deploying

In this section, the iApp gathers information about which Client Access services you are deploying. Some questions only appear depending on your answers to previous questions. These contingencies are noted at the beginning of the question description.

1. Do you want to customize the server pool settings?

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

Use settings recommended by F5

If you don't have a specific reason to customize the pool settings, leave this question set to this setting and continue with #2.

· Customize pool settings

If you need to modify individual pool options, select Customize pool settings and answer the following options that appear:

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

Select the load balancing method you want to use. We recommend the default, **Least Connections (member)**. See the BIG-IP documentation for a description of each method. If you chose a node-based load balancing method, such as Ratio (Node), and use a Ratio or Connection Limit (both optional), you must see <u>Adding Ratio or Connection Limit information to the nodes if using a node-based load balancing method on page 58 after completing the template.</u>

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

Select whether you want to enable Priority Group Activation to send traffic first to groups of servers you specify. The BIG-IP system load balances traffic according to the priority number you assign to each server.

Do not use Priority Group Activation

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

Use Priority Group Activation

Select this option if you want to enable Priority Group Activation. You will need to add a priority number in the Priority box to each server . A higher number indicates higher priority. Traffic is only sent to the servers with the highest priority, unless the number of available servers in that priority group falls below the value you specify as the minimum in the following question. The BIG-IP system then sends traffic to the group of servers with the next highest priority, and so on. See the BIG-IP documentation for more details.

a. What is the minimum number of active members in a group?

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

c. Do you want the BIG-IP system to queue TCP requests?

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

i Important TCP Request Queuing is an advanced feature and should be used only if you understand how it will affect your deployment, including application behavior and BIG-IP performance.

If you enable TCP Request Queuing, you must have a Connection Limit set on at least one of the nodes when configuring the Address/Port for the Client Access Server nodes.

• Do not queue TCP requests

Select this option if you do not want the BIG-IP system to queue TCP requests.

Queue TCP requests

Select this option if you want to enable TCP request queuing on the BIG-IP system.

a. What is the maximum number of TCP requests for the queue? Type the maximum number of requests you want to queue. We do not recommend using 0, which means unlimited and is only constrained by available memory.

b. How many milliseconds should requests remain in the queue? Type a number of milliseconds for the TCP request timeout value.

2. What IP address do you want to use for your virtual servers?

This question appears only if you selected Single IP address for all CAS connections in the previous section.

Specify a valid IP address to use for the BIG-IP virtual server. This virtual server address is used for all Client Access services. The BIG-IP system intelligently directs traffic to the appropriate service using an iRule created by the template.

3. Do you want to add any iRules to this combined virtual server?

If you chose to customize pool settings, you have the option of adding existing iRules to the virtual server. iRules allow an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. For more information on iRules, see https://devcentral.f5.com/HotTopics/iRules/tabid/1082202/Default.aspx.

i Important Improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommended you verify the impact of an iRule prior to deployment in a production environment.

If you want to add iRules, from the Options box, select the iRule(s) you want to include, and then click the Add (<<) button.

4. Are you deploying Outlook Web App (includes ECP)?

Select whether you are deploying Outlook Web App at this time. This includes the Exchange Control Panel (ECP).

Select this option if you are not deploying OWA at this time. You can always reconfigure the template later to add OWA.

• Yes

Select this option if you are deploying OWA at this time.

a. Which type of authentication do Outlook Web App clients use?

This question only appears if you selected to use BIG-IP APM and for the iApp to create a new Access profile

Choose whether your outlook Web App clients are using Forms-based authentication or Smart Card authentication. You must be using BIG-IP APM version 11.3 or later for Smart Card authentication support for OWA.

Outlook Web App clients use Forms-based authentication

Select this option if your Outlook Web App clients are using Forms-based authentication.

a. Would you like to display the OWA computer type and light version options on the APM logon page? Choose whether you want to display the computer type (public/shared vs private) and light version (Use the light version of Outlook Web App) options for OWA on the APM logon page.

No, do not display the OWA logon options

Select this option if you do not want to display the OWA logon options on the APM logon page.

Yes, display the OWA logon options

Select this option if you want users to see the OWA logon options on the BIG-IP APM logon page.

Note that in Exchange 2013, you must enable the logon page options by running a specific PowerShell command in the Exchange Management Shell prior to logging into OWA. See Powershell command for enabling the OWA logon options on page 87.

Note: For the blind and low vision experience to function correctly when accessing OWA with Internet Explorer 11, the OWA site must be added to the Compatibility View websites list. Consult Microsoft documentation for more information.

Outlook Web App clients use Smart Card authentication

Select this option if your OWA clients use Smart Card authentication and you are using BIG-IP APM v11.3 or later.

a. Specify a certificate from a CA that is trusted by this BIG-IP system for client-side processing of smart card authentication

Select the certificate you imported onto the BIG-IP system that is from a Certificate Authority and is trusted by the BIG-IP system for client-side processing of smart card authentication. This certificate must already be imported onto the system before you can select it.

b. Should BIG-IP APM restrict EAC access to members of the Exchange Organization Management Security Group? Exchange 2013 only

This question only appears if you selected Exchange 2013 as your version of Exchange and selected to provide secure authentication with BIG-IP APM.

Select whether you want the BIG-IP APM to restrict Exchange Administration Center (EAC) access to members of Exchange 2013's Organizational Management group. The BIG-IP APM module queries Active Directory group membership for the user making the request to EAC. If the user is not a member of the Organization Management group, the BIG-IP APM policy denies access.

No, do not restrict EAC access by group membership

Select this option and the BIG-IP APM will not restrict access to the EAC by group membership.

Yes, restrict EAC access by group membership

Select this option if you want to restrict EAC access to the Organization Management group. This adds an additional layer of security to your Exchange deployment, as the system denies access to the EAC from anyone who is not a member of the Organization Management group.

c. What IP address do you want to use for the OWA virtual server?

This question only appears if you selected Different IP addresses for different services in the previous section.

Specify the IP address the system will use for the Outlook Web App virtual server. Clients will use this IP address to access Outlook Web App.

d. Do you want to add custom iRules to this virtual server?

This question only appears if you selected Different IP addresses for different services in the previous section.

If you chose to customize pool settings, you have the option of adding existing iRules to this OWA virtual server. iRules allow an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. For more information on iRules, see https://devcentral.f5.com/HotTopics/iRules/tabid/1082202/Default.aspx.

If you want to add iRules, from the **Options** box, select the iRule(s) you want to include, and then click the Add (<<) button. Improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommended you verify the impact of an iRule prior to deployment in a production environment.

e. What are the IP addresses of your OWA servers?

This question only appears if you selected Each service will be handled by a unique set of Client Access Servers in the previous section.

Specify the IP addresses of your Outlook Web App servers. Click Add to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

5. Are you deploying Outlook Anywhere, EWS and OAB (or EWS only)?

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

• No, not deploying Outlook Anywhere, EWS, or OAB

Select this option if you are not deploying Outlook Anywhere at this time. You can always reconfigure the template at another time to add Outlook Anywhere to the configuration.

Yes, deploying EWS only

Select this option if you are <u>only</u> deploying Exchange Web Services at this time, and not Outlook Anywhere or Offline Address Book. In this case, the BIG-IP system sends any Offline Address Book traffic to the Exchange Web Services pool.

a. What IP address do you want to use for the Exchange Web Services virtual server?

https://devcentral.f5.com/HotTopics/iRules/tabid/1082202/Default.aspx.

This question only appears if you selected **Different IP addresses for different services** in the previous section. Specify the IP address the system will use for the Exchange Web Services virtual server.

b. Do you want to add custom iRules to this virtual server?

This question only appears if you selected **Different IP addresses for different services** in the previous section.

If you chose to customize pool settings, you have the option of adding existing iRules to this Exchange Web Services virtual server. iRules allow an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. For more information on iRules, see

If you want to add iRules, from the **Options** box, select the iRule(s) you want to include, and then click the Add (<<) button. Improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommended you verify the impact of an iRule prior to deployment in a production environment.

c. What are the IP addresses of your EWS servers?

This question only appears if you selected **Each service will be handled by a unique set of Client Access Servers** in the previous section.

Specify the IP addresses of your Exchange Web Services servers. Click **Add** to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

Yes, deploying Outlook Anywhere, EWS, and OAB

Select this option if you are deploying Outlook Anywhere, EWS, and OAB at this time.

i Important In Microsoft Exchange 2010, you must enable Outlook Anywhere on each of your Exchange Client Access

Servers before that service will be available. Outlook Anywhere is not enabled by default on Exchange Client Access Servers. See the Microsoft documentation for specific instructions. To prevent internal users from receiving a password prompt, your internal DNS must not have an 'A' record for the FQDN for Outlook Anywhere. This only applies if you are using Exchange 2010, are using RPC MAPI internally and Outlook Anywhere externally, and your internal clients do not have a route to the external Outlook Anywhere/EWS virtual server(s).

a. What IP address do you want to use for the Outlook Anywhere virtual server?

This question only appears if you selected **Different IP addresses for different services** in the previous section.

Specify the IP address the system will use for the Outlook Anywhere virtual server.

b. Do you want to add custom iRules to this virtual server?

This question only appears if you selected Different IP addresses for different services in the previous section.

If you chose to customize pool settings, you have the option of adding existing iRules to this Outlook Anywhere virtual server. iRules allow an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. For more information on iRules, see https://devcentral.f5.com/HotTopics/iRules/tabid/1082202/Default.aspx.

If you want to add iRules, from the **Options** box, select the iRule(s) you want to include, and then click the Add (<<) button. Improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommended you verify the impact of an iRule prior to deployment in a production environment.

c. What are the IP addresses of your Outlook Anywhere servers?

This question only appears if you selected **Each service will be handled by a unique set of Client Access Servers** in the previous section.

Specify the IP addresses of your Outlook Anywhere servers. Click **Add** to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

d. Which type of authentication do Outlook Anywhere clients use?

This question only appears if you chose to deploy BIG-IP APM and are using BIG-IP version 11.3 or later.

Choose whether your Outlook Anywhere clients use Basic or NTLM authentication. Beginning in BIG-IP version 11.3, the iApp supports using NTLM authentication for Outlook Anywhere.

• Outlook Anywhere clients use Basic Authentication

Select this option if your Outlook Anywhere clients use Basic Authentication. No further information is required, and you can continue with #6.

Outlook Anywhere clients use NTLM authentication

Select this option if your Outlook Anywhere clients use NTLM information. You must answer the following questions about your Active Directory implementation. Also see <u>Appendix E: Active Directory and Exchange Server configuration for NTLM on page 115</u> for important information and modifications for NTLM.

(i) Important Before completing this section, you must create a user account in the same domain that has been properly configured for Kerberos delegation. You must create an NTLM Machine Account object on the BIG-IP system to join this system to the Active Directory domain. See Creating an NTLM Machine Account on page 66.

a. Which NTLM machine account should be used for Kerberos delegation?

Select the NTLM Machine Account you created to join the BIG-IP system to the Active Directory domain. If you have not already created an NTLM Machine Account on the BIG-IP system, see <u>Creating an NTLM Machine Account on page 66</u>. You must either exit the template now and start over once you have created the NTLM Machine Account, or choose Outlook Anywhere Clients use Basic Authentication from the previous question, and then re-enter the template at a later time.

b. What is the Kerberos Key Distribution Center IP or FQDN?

Specify the IP address or fully qualified domain name of the Kerberos Key Distribution Center (KDC). If you type an FQDN, the BIG-IP system must be able to resolve the address. Otherwise, use the IP address.

c. What is the name of the Kerberos Realm?

Specify the name of the Kerberos Realm. While this name should be in all capital letters, the iApp automatically turns any lower case letters to capital.

d. What is the user name for the Active Directory delegation account you created?

Specify the user name for the Active Directory delegation account you created. This account must be correctly configured in Active Directory for Kerberos delegation. See <u>Appendix E: Active Directory and Exchange Server configuration for NTLM on page 115</u> details.

e. What is the associated password?

Specify the password associated with the account.

f. How do you want to construct the Kerberos ticket request?

Select whether you want to use DNS reverse lookups or the Outlook Anywhere Host header to construct the ticket request.

Use DNS reverse lookups

Select this option to use DNS reverse lookups to build the Kerberos ticket request. Note that you must configure a reverse lookup zone containing a PTR record for each Client Access Server on a DNS server that is accessible from this BIG-IP system. Consult your DNS documentation for specific instructions.

Use the Outlook Anywhere host header

Select this option to use the Outlook Anywhere Host header to construct the ticket request. To use the host header value, you must configure IIS Application Pools for Outlook Anywhere, Autodiscover, and Exchange Web Services to run using the previously created Active Directory user account for Kerberos delegation. See *Appendix E: Active Directory and Exchange Server configuration for NTLM on page 115*.

6. Are you deploying ActiveSync?

Select whether you are deploying ActiveSync at this time.

No

Select this option if you are not deploying ActiveSync at this time. You can always reconfigure the template at another time to add ActiveSync to the configuration.

Yes

Select this option if you are deploying ActiveSync at this time. See iPhones and other iOS devices are displaying invalid certificate messages after deploying the iApp for ActiveSync on page 61 for important information.

a. What IP address do you want to use for the ActiveSync virtual server?

This question only appears if you selected Different IP addresses for different services in the previous section.

Specify the IP address the system will use for the ActiveSync virtual server.

b. Do you want to add custom iRules to this virtual server?

This question only appears if you selected Different IP addresses for different services in the previous section.

If you chose to customize pool settings, you have the option of adding existing iRules to this ActiveSync virtual server. iRules allow an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. For information on iRules, see https://devcentral.f5.com/HotTopics/iRules/tabid/1082202/Default.aspx.

If you want to add iRules, from the **Options** box, select the iRule(s) you want to include, and then click the Add (<<) button. Improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommended you verify the impact of an iRule prior to deployment in a production environment.

c. What are the IP addresses of your ActiveSync servers?

This question only appears if you selected Each service will be handled by a unique set of Client Access Servers in the previous section.

Specify the IP addresses of your Outlook Anywhere servers. Click Add to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

7. Are you deploying Autodiscover?

Select whether you are deploying Autodiscover at this time.

No

Select this option if you are not deploying Autodiscover at this time. You can always reconfigure the template at another time to add Autodiscover to the configuration.

Yes

Select this option if you are deploying Autodiscover at this time.



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

a. What IP address do you want to use for the Autodiscover virtual server?

This question only appears if you selected Different IP addresses for different services in the previous section.

Specify the IP address the system will use for the Autodiscover virtual server.

b. Do you want to add custom iRules to this virtual server?

This question only appears if you selected Different IP addresses for different services in the previous section.

If you chose to customize pool settings, you have the option of adding existing iRules to this Autodiscover virtual server. iRules allow an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. For iRule information, see https://devcentral.f5.com/HotTopics/iRules/tabid/1082202/Default.aspx.

If you want to add iRules, from the Options box, select the iRule(s) you want to include, and then click the Add (<<) button. Improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommended you verify the impact of an iRule prior to deployment in a production environment.

c. What are the IP addresses of your Autodiscover servers?

This question only appears if you selected Each service will be handled by a unique set of Client Access Servers in the previous section.

Specify the IP addresses of your Autodiscover servers. Click Add to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

8. Are you deploying RPC Client Access (MAPI)? Exchange 2010 only

This question does not appear if you are deploying the template for Exchange 2013. Exchange Server 2013 Client Access Servers do not offer MAPI as a connection option.

Select whether you are deploying RPC Client Access (MAPI) for your Exchange 2010 deployment at this time.

No

Select this option if you are not deploying RPC Client Access at this time. You can always reconfigure the template at another time to add it to the configuration.

Yes

Select this option if you are deploying RPC Client Access at this time.



Warning You must enable and configure a Client Access Array in your Exchange Server site before RPC Client Access will function. See Creating a new Client Access Array on page 114 for more information. If deploying RPC Client Access, you must also deploy Outlook Anywhere, to properly handle EWS (Exchange Web Services) traffic.

a. Use the default dynamic range of ports for RPC Client Access traffic or set static ports?

Select whether you want to use the default dynamic range of ports for RPC Client Access, or if you have configured your Client Access servers to use specific ports outside the default range.

· Use the default dynamic port range

Select this option to configure the iApp to use the default port range. If you choose the default dynamic range of ports, no additional information is necessary, continue with the next question.

Set static ports

Select this option if you want to set static ports for RPC Client Access.

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

a. Which port will you use for MAPI?

Specify the port you want to set for MAPI.

b. Which port will you use for the Address Book?

Specify the port you want to use for the Address book.

b. What IP address do you want to use for the RPC Client Access virtual server?

This question only appears if you selected Different IP addresses for different services in the previous section.

Specify the IP address the system will use for the RPC Client Access virtual server.

c. What are the IP addresses of your RPC Client Access servers?

This question only appears if you selected Each service will be handled by a unique set of Client Access Servers in

Specify the IP addresses of your RPC Client Access servers. Click Add to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

9. Are you deploying POP3?

Select whether you are deploying POP3 at this time.

Select this option if you are not deploying POP3 at this time. You can always reconfigure the template at another time to add POP3 to the configuration.

Yes

Select this option if you are deploying POP3 at this time.

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

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

a. What IP address do you want to use for the POP3 virtual server?

This question only appears if you selected Different IP addresses for different services in the previous section.

Specify the IP address the system will use for the POP3 virtual server.

b. What are the IP addresses of your POP3 servers?

This question only appears if you selected **Each service will be handled by a unique set of Client Access Servers** in the previous section.

Specify the IP addresses of your POP3 servers. Click **Add** to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

10. Are you deploying IMAP4?

Select whether you are deploying IMAP4 at this time.

No

Select this option if you are not deploying IMAP4 at this time. You can always reconfigure the template at another time to add IMAP4 to the configuration.

Yes

Select this option if you are deploying IMAP4 at this time.

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

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

c. What IP address do you want to use for the IMAP4 virtual server?

This question only appears if you selected Different IP addresses for different services in the previous section.

Specify the IP address the system will use for the IMAP4 virtual server.

d. What are the IP addresses of your IMAP4 servers?

This question only appears if you selected **Each service will be handled by a unique set of Client Access Servers** in the previous section.

Specify the IP addresses of your IMAP4 servers. Click **Add** to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

11. What are the IP Addresses of your Client Access Servers?

This question only appears if you selected **All services will be handled by a unique set of Client Access Servers** in the previous section.

If you chose that each Client Access service will be handled by the same Client Access Servers, the iApp asks for the IP addresses of the Client Access Servers. Type the IP addresses. Click the **Add** button to include additional servers.

If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

Server Health Monitors

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

1. Do you want to use advanced or simple server health monitors?

Choose whether you want to use advanced or simple health monitors to check the availability of the Client Access Servers:

· Use simple monitors

Simple monitors check network connectivity but do not perform actual logins. If you use simple monitors, the BIG-IP LTM may not be able to completely determine status of Client Access services. In this case, the monitor interval is set to 10 seconds automatically, no matter what number is in the previous question.

· Use advanced monitors

If you choose advanced monitors, the BIG-IP system performs logins to most of the Client Access services (all except RPC/MAPI in Exchange 2010 and Forms-based Outlook Web App) and checks for valid content in the response. Because these monitors attempt to access a specific mailbox, they can more accurately determine the actual health of Client Access services. However, account maintenance and Mailbox status must become a part of your monitoring strategy. For example, if an account used for monitoring is locked or deleted, the monitor will mark the services down for all users.

i Important F5's advanced monitors for Autodiscover, Exchange Web Services, and Outlook Anywhere support Basic and NTLMv1 authentication only. See Advanced monitors for Autodiscover, EWS, and Outlook Anywhere only support Basic and NTLMv1 authentication on page 60 for more information.

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

a. What email address do you want to use for the advanced monitors?

This question only appears if you specified you are deploying Autodiscover and/or Exchange Web Services .

Type the email address associated with the account you are going to monitor (and that you specify in the following question).

b. Which mailbox account should be used for the monitors?

Type a mailbox account for use in the advanced monitors. This name corresponds to the account name field in Active Directory (rather than the email address).

c. What is the password for that mailbox account?

Type the associated password.

d. What is the domain name of the user account for the monitors?

Type the Domain name for the user account. This domain can be entered in either FQDN (mydomain.example.com) or NetBIOS (MYDOMAIN) format.

e. Do you want to monitor a second mailbox?

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

• Monitor only one mailbox

Select this option if you do not want the BIG-IP system to monitor a second mailbox. Continue with #3.

• Monitor a second mailbox (recommended)

Select this option if you want the BIG-IP system to monitor a second mailbox. You must answer the following:

a. Which email address do you want to use for the second advanced monitor?

This question only appears if you specified you are deploying Autodiscover and/or Exchange Web Services .

Type the email address associated with the account you are going to monitor (and that you specify in the following question).

b. Which mailbox account should be used for the second monitor?

Type a mailbox account for use in the second monitors. Again, this name corresponds to the account name field in Active Directory (rather than the email address).

c. What is the password for that mailbox account?

Type the associated password.

d. What is the domain name of the user account for the second monitors?

Type the Domain name for the second user account. This domain can be entered in either FQDN (mydomain. example.com) or NetBIOS (MYDOMAIN) format.

2. Which authentication method have you configured for OWA?

This question only appears if you specified you are deploying Outlook Web App.

If you configured the iApp to deploy Outlook Web App at this time, choose the authentication method you have configured for Outlook Web App. The health monitors will be customized to accommodate the authentication method you are using.

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

OWA uses the default Forms-Based authentication

Select this option if you are using Forms-based authentication.

If you chose Forms-based authentication, the BIG-IP system does not perform an actual login to the service, but checks the availability of the forms-based authentication page.

· OWA uses Basic or Windows Integrated authentication

Select this option if you are using Basic/Windows Integrated authentication.

3. How many seconds should pass between health checks?

Specifies how often the system checks the health of the servers. We recommend the default of 30 seconds. The maximum value for the interval is 28,799 seconds.

4. Are you using the same FQDN for all HTTP-based services?

This question only appears if you specified you are using a <u>single IP address</u> for all CAS connections. If you selected Different IP addresses for different services, continue with #5.

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

· One FQDN for all HTTP services

Select this option if you are using a single FQDN for all HTTP-based Client Access services.

a. What is the FQDN that you use for your HTTP-based CAS services?

Specify the fully qualified domain name you are using for all of the HTTP-based CAS services.

· Different FQDNs for each HTTP service

Select this option if you are using separate FQDNs for each HTTP-based CAS service. Additional questions appear. When you are finished adding the FQDNs, continue with *Additional Steps*.

a. What FQDN do you use for the OWA service?

This question only appears if you specified you are deploying Outlook Web App. Specify the fully qualified domain name you use for your Outlook Web App service.

b. What FQDN do you use for the Outlook Anywhere service?

This question only appears if you specified you are deploying Outlook Anywhere. Specify the fully qualified domain name you use for your Outlook Anywhere service.

c. What FQDN do you use for the ActiveSync service?

This question only appears if you specified you are deploying ActiveSync. Specify the fully qualified domain name you use for your ActiveSync service.

d. What FQDN do you use for the Autodiscover service?

This question only appears if you specified you are deploying Autodiscover. Specify the fully qualified domain name you use for your Autodiscover service.

5. What FQDN do you use for the OWA service?

This question only appears if you specified you are using different IP addresses for different services.

Specify the fully qualified domain name you use for your Outlook Web App service.

6. What FQDN do you use for the Outlook Anywhere service?

This question only appears if you specified you are using <u>different IP addresses</u> for different services.

Specify the fully qualified domain name you use for your Outlook Anywhere service.

7. What FQDN do you use for the ActiveSync service?

This question only appears if you specified you are using <u>different IP addresses</u> for different services.

Specify the fully qualified domain name you use for your ActiveSync service.

8. What FQDN do you use for the Autodiscover service?

This question only appears if you specified you are using different IP addresses for different services.

Specify the fully qualified domain name you use for your Autodiscover service.

Additional Steps

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

Finished

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

Continue with Next steps on page 57.

Configuring the LTM to receive HTTP-based Client Access traffic forwarded by an APM

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

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

Analytics

This section of the template asks questions about Analytics. The Application Visibility Reporting (AVR) module allows you to view statistics specific to your Microsoft Exchange implementation. AVR is available on all BIG-IP systems v11 and later, however you must have the AVR provisioned for this option to appear. Note that this is only for application visibility reporting, you can view object-level statistics from the BIG-IP without provisioning AVR.

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

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

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

Choose whether you want to enable AVR for Analytics.

• No, do not enable Analytics

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

• Yes, enable Analytics using AVR

If you choose to enable Analytics, select Yes from the list, and then answer the following questions.

a. Use the default Analytics profile or select a custom profile?

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

· Select a custom Analytics profile

Select this option if you have already created a custom Analytics profile for Exchange Server.

a. Which Analytics profile do you want to use?

From the list, select the appropriate Analytics profile.

Use default profile

Select this option if you have not yet created a custom Analytics profile for Microsoft Exchange. We do not recommend using the default profile.

Application Firewall Manager (BIG-IP AFM)

This section gathers information about BIG-IP Advanced Firewall Manager if you want to use it to protect the Exchange deployment. For more information on configuring AFM, see http://support.f5.com/kb/en-us/products/big-ip-afm.html, and then select your version.

1. Do you want to use BIG-IP AFM to protect your application?

Choose whether you want to use BIG-IP AFM, F5's network firewall, to secure this Exchange deployment. If you choose to use BIG-IP AFM, you can restrict access to the Exchange virtual server(s) to a specific network or IP address. See the BIG-IP AFM documentation for specific details on configuring AFM.

No, do not use Application Firewall Manager

Select this option if you do not want to enable BIG-IP AFM at this time. You can always re-enter the template at a later date to enable BIG-IP AFM. Continue with the next section.

Select an existing AFM policy from the list

If you already created a BIG-IP AFM policy for this implementation, select it from the list. Continue with c.

▶ Yes, use F5's recommended AFM configuration

Select this option if you want to enable BIG-IP AFM using F5's recommended configuration.

a. Do you want to restrict access to your application by network or IP address?

Choose whether you want to restrict access to the Exchange implementation via the BIG-IP virtual server.

No, do not restrict source addresses (allow all sources)

By default, the iApp configures the AFM to accept traffic destined for the Exchange virtual server(s) from all sources. If you do not have a need to restrict access to the virtual server, leave this option selected and then continue with \boldsymbol{b} .

Restrict source addresses

Select this option if you want to restrict access to the Exchange virtual server(s) by IP address or network address.

i). What IP or network addresses should be allowed to access your application?

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

b. How do you want to control access to your application from sources with a low reputation score?

The BIG-IP AFM uses an IP intelligence database to categorize IP addresses coming into the system. Choose what you want the system to do for sources that are attempting to access the Exchange virtual server(s) with a low reputation score. For more information, see the BIG-IP AFM documentation.

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

Allow all sources regardless of reputation

Select this option to allow all sources, without taking into consideration the reputation score.

▶ Reject access from sources with a low reputation

Select this option to reject access to the Exchange virtual server(s) from any source with a low reputation score.

▶ Allow but log access from sources with a low reputation

Select this option to allow access to the Exchange virtual server(s) from sources with a low reputation score, but add an entry for it in the logs.

c. Would you like to stage a policy for testing purposes?

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

Do not apply a staging policy

Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.

Select an existing policy from the list

If you have already created a firewall policy for this implementation, select it from the list. Only policies that already exist on the system appear in the list. To create a new policy, on the Main tab, click **Security > Network Firewall > Policies**. Specific instructions for creating a firewall policy is outside the scope of this iApp and deployment guide.

d. Which logging profile would you like to use?

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

Do not apply a logging profile

Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.

Select an existing logging profile from the list

If you have already created a logging profile for this implementation, select it from the list. You must create a profile before it is available in the list. To create a logging profile, on the Main tab, click **Security > Event Logs** > **Logging Profiles**. Specific instructions for creating a logging profile is outside the scope of this iApp and deployment guide. See the online help or the *About Local Logging with the Network Firewall* chapter of the *BIG-IP Network Firewall: Policies and Implementations* guide for more information.

Tell us about your deployment

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

1. Will incoming traffic arrive at this BIG-IP system encrypted or unencrypted?

Select whether any of the HTTP-based, POP3 and IMAP4 traffic will be encrypted or not when it arrives on this system. Because you may have configured to offload SSL/TLS traffic at the BIG-IP APM that is sending Client Access traffic to this device, the traffic may be arriving at this system unencrypted.

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

Encrypted

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

a. Do you want to re-encrypt this traffic to your Client Access Servers?

If you are using Exchange 2010 or Exchange 2013 SP1 and later, and want the BIG-IP system to offload SSL processing from the Client Access Servers, select **Do not re-encrypt (SSL Offload)** from the list. Offloading SSL on the BIG-IP system can extend Exchange Server server capacity.

(i) Important If you are deploying Exchange Server 2013 and have not installed SP1 or later, you must choose Re-encrypt (SSL Bridging).

Do not re-encrypt (SSL Offload)

Select this option if you want to offload SSL processing onto the BIG-IP system. If you choose SSL Offload, you must have followed the instructions described in the prerequisites for configuring the Exchange Server: http://social.technet.microsoft.com/wiki/contents/articles/how-to-configure-ssl-offloading-in-exchange-2010.aspx.

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

The iApp can create a new Client SSL profile, or if you have created a Client SSL profile which contains the appropriate SSL certificate and key for your Exchange implementation, you can select it from the list.

- Select the Client SSL profile you created from the list
 If you manually created a Client SSL profile, select it from the list, and then continue with #2.
- Create a new Client SSL profile

Select this option if you want the iApp to create a new Client SSL profile.

a. Which SSL certificate do you want to use?

Select the SSL certificate you imported onto the BIG-IP system for decrypting client connections. If you have not yet imported a certificate, you can leave the default selections and reconfigure this iApp after obtaining the certificates. The deployment will not function correctly until you have selected the correct certificates here.

Note: Any certificate that you obtain with multiple names must be in SAN (Subject Alternative Name) format, not SNI (Server Name Indication) format. For more information on SAN certificates, see Subject Alternative Name (SAN) SSL Certificates on page 113.

b. Which SSL key do you want to use?

Select the associated key from the list.

Re-encrypt (SSL Bridging)

Select this option if your implementation requires encrypted traffic to the Client Access Servers, or you are using Exchange 2013. The BIG-IP system unencrypts, and then re-encrypts the traffic headed for the Client Access Servers.

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

The iApp can create a new Client SSL profile, or if you have created a Client SSL profile which contains the appropriate SSL certificate and key for your Exchange implementation, you can select it from the list.

Select the Client SSL profile you created from the list If you manually created a Client SSL profile, select it from the list, and then continue with #2.

Create a new Client SSL profile

Select this option if you want the iApp to create a new Client SSL profile.

a. Which SSL certificate do you want to use?

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

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

Note: Any certificate that you obtain with multiple names must be in SAN (Subject Alternative Name) format, not SNI (Server Name Indication) format. For more information on SAN certificates, see Subject Alternative Name (SAN) SSL Certificates on page 113.

b. Which SSL key do you want to use?

Select the associated key from the list.

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

Select whether you want the iApp to create an F5 recommended Server SSL profile, or if you want to choose a Server SSL profile you already created.

(i) Important If you are configuring SSL Bridging and using BIG-IP version 11.4.x, you must see When using SSL Bridging and BIG-IP version 11.4.x, pool members may be marked down or you may experience connection resets and TLS errors logged to the Client Access servers on page 62.

Select the Server SSL profile you created from the list

If you have previously created a Server SSL profile for your Exchange implementation, from the list, select the existing Server SSL profile you created.

Create a new Server SSL profile

Select this option if you want the iApp to create a new Server SSL profile.

The default, F5 recommended Server SSL profile uses the serverssl parent profile. For information about the ciphers used in the Server SSL profile, see

http://support.f5.com/kb/en-us/solutions/public/8000/800/sol8802.html.

Unencrypted

Select this option if Client Access traffic is arriving at this BIG-IP system unencrypted (typically because you configured to offload SSL/TLS traffic at the BIG-IP APM that is sending Client Access traffic to this device).

a. Do you want to encrypt the traffic to your Client Access Servers?

If you are using Exchange 2010 or Exchange 2013 SP1 and later, and want the BIG-IP system to offload SSL processing from the Client Access Servers, select Do not encrypt (SSL Offload) from the list. Offloading SSL on the BIG-IP system can extend Exchange Server server capacity.

(i) Important If you are deploying Exchange Server 2013 and have not installed SP1 or later, you must choose Re-encrypt (SSL Bridging).

Do not encrypt (SSL Offload)

Select this option if you do not want the BIG-IP system to encrypt the traffic destined for the Client Access servers. The BIG-IP system does not modify the traffic, and you can continue with the next question.

• Encrypt (SSL Bridging)

Select this option if your implementation requires encrypted traffic to the Client Access Servers, or you are using Exchange 2013 prior to SP1. The BIG-IP system encrypts the traffic headed for the Client Access Servers.

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

Select whether you want the iApp to create an F5 recommended Server SSL profile, or if you want to choose a Server SSL profile you already created.

(i) Important If you are configuring SSL Bridging and using BIG-IP version 11.4.x, you must see When using SSL Bridging and BIG-IP version 11.4.x, pool members may be marked down or you may experience connection resets and TLS errors logged to the Client Access servers on page 62.

• Select the Server SSL profile you created from the list

If you have previously created a Server SSL profile for your Exchange implementation, from the list, select the existing Server SSL profile you created.

· Create a new Server SSL profile

Select this option if you want the iApp to create a new Server SSL profile.

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

2. Where will your BIG-IP virtual servers be in relation to your Client Access Servers?

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

· Same subnet for BIG-IP virtual servers and Client Access Servers

Select this option if the BIG-IP virtual servers and the Client Access Servers are on the same subnet. In this case SNAT is configured on the BIG-IP virtual server and you must specify the number of concurrent connections.

a. What is the maximum number of concurrent users you expect per Client Access Server?

Select whether you expect more or fewer than 6,000 concurrent users to each Client Access Server. This answer is used to determine what type of SNAT that system uses. A SNAT is an object that maps the source client IP address in a request to a translation address defined on the BIG-IP device.

Note: For specific information on SNAT Pools, including why we chose 6,000 concurrent users per Client Access Server, see Maximum number of concurrent users: SNAT Pool guidance on page 113.

Fewer than 6000

Select this option if you expect fewer than 6000 concurrent users per Client Access Server. With this option, the system applies SNAT Auto Map, which doesn't require any additional IP addresses, as the system uses an existing self IP address for translation.

More than 6000

Select this option if you expect more than 6,000 users at one time to each server. With this option, the iApp creates a SNAT Pool (or you can choose one you created), for which you need one IP address for each 6,000 users you expect.

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

Select whether you want the system to create a new SNAT Pool, or if you have already created a SNAT pool for this implementation.

Select the SNAT pool you created from the list

If you have previously created a SNAT Pool for your Exchange implementation, select it from the list.

Create a new SNAT pool

If you have not created a custom SNAT pool, select this option for the iApp to create a new one.

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

Specify one otherwise unused IP address for every 6,000 concurrent connections, or fraction thereof. Click **Add** for additional rows.

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

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

a. How have you configured routing on your Client Access Servers?

Select whether the Client Access Servers use this BIG-IP system's Self IP address as their default gateway.

Client Access Servers do NOT use BIG-IP as their default gateway

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

a. What is the maximum number of concurrent users you expect per Client Access Server?

Select whether you expect more or fewer than 6,000 concurrent users to each Client Access Server. This answer is used to determine what type of SNAT that system uses. A SNAT is an object that maps the source client IP address in a request to a translation address defined on the BIG-IP device.

Note: For specific information on SNAT Pools, including why we chose 6,000 concurrent users per Client Access Server, see Maximum number of concurrent users: SNAT Pool guidance on page 113.

Fewer than 6000

Select this option if you expect fewer than 6000 concurrent users per Client Access Server. With this option, the system applies SNAT Auto Map, which doesn't require any additional IP addresses, as the system uses an existing self IP address for translation.

More than 6000

Select this option if you expect more than 6,000 users at one time to each server. With this option, the iApp creates a SNAT Pool (or you can choose one you created), for which you need one IP address for each 6,000 users you expect.

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

Select whether you want the system to create a new SNAT Pool, or if you have already created a SNAT pool for this implementation.

- Select the SNAT pool you created from the list
 If you have previously created a SNAT Pool for your Exchange implementation, select it from the list.
- Create a new SNAT pool

If you have not created a custom SNAT pool, select this option for the iApp to create a new one.

- a. Which IP addresses do you want to use for the SNAT pool?
 Specify one otherwise unused IP address for every 6,000 concurrent connections, or fraction thereof. Click Add for additional rows.
 - i Important If you choose more than 6,000 users, but do not specify enough SNAT pool addresses, after the maximum connection limit of 6,000 concurrent users per server is reached, new requests fail.
- Client Access Servers use the BIG-IP as their default gateway
 Select this option if the Client Access Servers use the BIG-IP system as their default gateway. In this scenario, no additional configuration is necessary to ensure correct server response handling.

3. How are you distributing the CAS protocols between servers?

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

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

- All services will be handled by the same set of Client Access Servers
 Choose this option if you are using the same Client Access Servers for all of your Exchange Client Access services.
- Each service will be handled by a unique set of Client Access Servers

 Choose this option if you are using different sets of Client Access Servers for each Client Access service.

Tell us about which services you are deploying

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

1. Would you like to customize the server pool settings?

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

Use settings recommended by F5

If you don't have a specific reason to customize the pool settings, leave this question at Use settings recommended by F5 and continue with #2.

Customize pool settings

If you have need to modify individual pool options, select Customize pool settings and answer the following options:

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

Select the load balancing method you want to use. We recommend the default, Least Connections (member). See the BIG-IP documentation for a description of each method. If you chose a node-based load balancing method (such as Ratio (node)), and use a Ratio or Connection Limit (both optional), you must see Adding Ratio or Connection Limit information to the nodes if using a node-based load balancing method on page 58 after completing the template.

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

Select whether you want to enable Priority Group Activation to send traffic first to groups of servers you specify. The BIG-IP system load balances traffic according to the priority number you assign to each server.

Do not use Priority Group Activation

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

Use Priority Group Activation

Select this option if you want to enable Priority Group Activation. You will need to add a priority number in the Priority box to each server. A higher number indicates higher priority. Traffic is only sent to the servers with the highest priority, unless the number of available servers in that priority group falls below the value you specify as the minimum in the following question. The BIG-IP system then sends traffic to the group of servers with the next highest priority, and so on. See the BIG-IP documentation for more details.

a. What is the minimum number of active members in a group?

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

c. Do you want the BIG-IP system to queue TCP requests?

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

(i) Important TCP Request Queuing is an advanced feature and should be used only if you understand how it will affect your deployment, including application behavior and BIG-IP performance.

> If you enable TCP Request Queuing, you must have a Connection Limit set on at least one of the nodes when configuring the Address/Port for the Client Access Server nodes.

Do not queue TCP requests

Select this option if you do not want the BIG-IP system to queue TCP requests.

Queue TCP requests

Select this option if you want to enable TCP request queuing on the BIG-IP system.

a. What is the maximum number of TCP requests for the queue?

Type the maximum number of requests you want to queue. We do not recommend using 0, which means unlimited and is only constrained by available memory.

b. How many milliseconds should requests remain in the queue?

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

2. What IP address do you want to use for your virtual servers?

This question appears only if you selected Single IP address for all CAS connections in the previous section.

Specify a valid IP address to use for the BIG-IP virtual server. This virtual server address is used for all Client Access services. The BIG-IP system intelligently directs traffic to the appropriate service using an iRule created by the template.

3. Do you want to add any iRules to this combined virtual server?

If you chose to customize pool settings, you have the option of adding existing iRules to the virtual server. iRules allow an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. For more information on iRules, see https://devcentral.f5.com/HotTopics/iRules/tabid/1082202/Default.aspx.

i Important Improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommended you verify the impact of an iRule prior to deployment in a production environment.

If you want to add iRules, from the Options box, select the iRule(s) you want to include, and then click the Add (<<) button.

4. Are you deploying Outlook Web App (includes ECP)?

Select whether you are deploying Outlook Web App at this time. This includes the Exchange Control Panel (ECP).

No

Select this option if you are not deploying OWA at this time. You can reconfigure the template at another time to add OWA.

Yes

Select this option if you are deploying Outlook Web Access at this time.

a. What are the IP addresses of your OWA servers?

This question only appears if you selected **Each service will be handled by a unique set of Client Access Servers** in the previous section.

Specify the IP addresses of your Outlook Web App servers. Click **Add** to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

5. Are you deploying Outlook Anywhere, EWS and OAB (or EWS only)?

Select whether you are deploying Outlook Anywhere, Exchange Web Services (EWS), Offline Address Book (OAB), or EWS only.

· No, not deploying Outlook Anywhere, EWS, or OAB

Select this option if you are not deploying Outlook Anywhere at this time. You can always reconfigure the template at another time to add Outlook Anywhere to the configuration.

Yes, deploying EWS only

Select this option if you are <u>only</u> deploying Exchange Web Services at this time, and not Outlook Anywhere or Offline Address Book. In this case, the BIG-IP system sends any Offline Address Book traffic to the Exchange Web Services pool.

a. What are the IP addresses of your EWS servers?

This question only appears if you selected **Each service will be handled by a unique set of Client Access Servers** in the previous section.

Specify the IP addresses of your Exchange Web Services servers. Click **Add** to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

· Yes, deploying Outlook Anywhere, EWS, and OAB

Select this option if you are deploying Outlook Anywhere, EWS, and OAB at this time.

i Important In Microsoft Exchange 2010, you must enable Outlook Anywhere on each of your Exchange Client Access Servers before that service will be available. Outlook Anywhere is not enabled by default on Exchange Client Access Servers. See the Microsoft documentation for specific instructions. To prevent internal users from receiving a password prompt, your internal DNS must not have an 'A' record for the FQDN for Outlook Anywhere. This only applies if you are using Exchange 2010, are using RPC MAPI internally and Outlook Anywhere externally, and your internal clients do not have a route to the external Outlook Anywhere/EWS virtual server(s).

a. What are the IP addresses of your Outlook Anywhere servers?

This question only appears if you selected **Each service will be handled by a unique set of Client Access Servers** in the previous section.

Specify the IP addresses of your Outlook Anywhere servers. Click **Add** to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

6. Are you deploying ActiveSync?

Select whether you are deploying ActiveSync at this time.

No

Select this option if you are not deploying ActiveSync at this time. You can always reconfigure the template at another time to add ActiveSync to the configuration.

Yes

Select this option if you are deploying ActiveSync at this time. Be sure to see iPhones and other iOS devices are displaying. invalid certificate messages after deploying the iApp for ActiveSync on page 61 for important information.

a. What are the IP addresses of your ActiveSync servers?

This question only appears if you selected Each service will be handled by a unique set of Client Access Servers in the previous section.

Specify the IP addresses of your Outlook Anywhere servers. Click Add to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

7. Are you deploying Autodiscover?

Select whether you are deploying Autodiscover at this time.

Select this option if you are not deploying Autodiscover at this time. You can always reconfigure the template at another time to add Autodiscover to the configuration.

Select this option if you are deploying Autodiscover at this time.



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

a. What are the IP addresses of your Autodiscover servers?

This question only appears if you selected Each service will be handled by a unique set of Client Access Servers in the previous section.

Specify the IP addresses of your Autodiscover servers. Click Add to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit.

8. Are you deploying RPC Client Access (MAPI)? Exchange 2010 only

This question does not appear if you are deploying the template for Exchange 2013. Exchange Server 2013 Client Access Servers do not offer MAPI as a connection option.

Select whether you are deploying RPC Client Access (MAPI) for your Exchange 2010 deployment at this time.

Nο

Select this option if you are not deploying RPC Client Access at this time. You can always reconfigure the template at another time to add it to the configuration.

Yes

Select this option if you are deploying RPC Client Access at this time.



/N Warning You must enable and configure a Client Access Array in your Exchange Server site before RPC Client Access will function. See Creating a new Client Access Array on page 114 for more information. If deploying RPC Client Access, you must also deploy Outlook Anywhere, to properly handle EWS (Exchange Web Services) traffic.

a. Use the default dynamic range of ports for RPC Client Access traffic or set static ports?

Select whether you want to use the default dynamic range of ports for RPC Client Access, or if you have configured your Client Access servers to use specific ports outside the default range.

• Use the default dynamic port range

Select this option to configure the iApp to use the default port range. If you choose the default dynamic range of ports, no additional information is necessary, continue with the next question.

· Set static ports

Select this option if you want to set static ports for RPC Client Access.

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

a. Which port will you use for MAPI?

Specify the port you want to set for MAPI.

b. Which port will you use for the Address Book?

Specify the port you want to use for the Address book.

b. What are the IP addresses of your RPC Client Access servers?

This question only appears if you selected **Each service will be handled by a unique set of Client Access Servers** in the previous section.

Specify the IP addresses of your RPC Client Access servers. Click **Add** to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

9. Are you deploying POP3?

Select whether you are deploying POP3 at this time.

No

Select this option if you are not deploying POP3 at this time. You can always reconfigure the template at another time to add POP3 to the configuration.

Yes

Select this option if you are deploying POP3 at this time.

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

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

a. What are the IP addresses of your POP3 servers?

This question only appears if you selected **Each service will be handled by a unique set of Client Access Servers** in the previous section.

Specify the IP addresses of your POP3 servers. Click **Add** to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

2. Are you deploying IMAP4?

Select whether you are deploying IMAP4 at this time.

No

Select this option if you are not deploying IMAP4 at this time. You can always reconfigure the template at another time to add IMAP4 to the configuration.

Yes

Select this option if you are deploying IMAP4 at this time.

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

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

a. What are the IP addresses of your IMAP4 servers?

This question only appears if you selected **Each service will be handled by a unique set of Client Access Servers** in the previous section.

Specify the IP addresses of your IMAP4 servers. Click **Add** to include additional servers. If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

2. What are the IP Addresses of your Client Access Servers?

This question only appears if you selected **All services will be handled by a unique set of Client Access Servers** in the previous section.

If you chose that each Client Access service will be handled by the same Client Access Servers, the iApp asks for the IP addresses of the Client Access Servers. Type the IP addresses. Click the **Add** button to include additional servers.

If you chose to have the BIG-IP system queue TCP requests, you must specify a Connection Limit. If you chose to enable Priority Group Activation, you must specify a Priority.

Server Health Monitors

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

1. Do you want to use advanced or simple server health monitors?

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

Use simple monitors

Simple monitors check network connectivity but do not perform actual logins. If you use simple monitors, the BIG-IP LTM may not be able to completely determine status of Client Access services. In this case, the monitor interval is set to 10 seconds automatically, no matter what number is in the previous question.

Use advanced monitors

If you choose advanced monitors, the BIG-IP system performs logins to most of the Client Access services (all except RPC/MAPI in Exchange 2010, and Forms-based Outlook Web App) and checks for valid content in the response. Because these monitors attempt to access a specific mailbox, they can more accurately determine the actual health of Client Access services. However, account maintenance and Mailbox status must become a part of your monitoring strategy. For example, if an account used for monitoring is locked or deleted, the monitor will mark the services down for all users.

(i) Important F5's advanced monitors for Autodiscover, Exchange Web Services, and Outlook Anywhere support Basic and NTLMv1 authentication only. See Advanced monitors for Autodiscover, EWS, and Outlook Anywhere only support Basic and NTLMv1 authentication on page 60.

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

a. What email address do you want to use for the advanced monitors?

This question only appears if you specified you are deploying Autodiscover and/or Exchange Web Services .

Type the email address associated with the account you are going to monitor (that you specify in the following question).

b. Which mailbox account should be used for the monitors?

Type a mailbox account for use in the advanced monitors. This name corresponds to the account name field in Active Directory (rather than the email address).

c. What is the password for that mailbox account?

Type the associated password. Note that credentials are stored in plain text on this BIG-IP system.

d. What is the domain name of the user account for the monitors?

Type the Domain name for the user account. This domain can be entered in either FQDN (mydomain.example.com) or NetBIOS (MYDOMAIN) format.

e. Do you want to monitor a second mailbox?

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

No

Select this option if you do not want the BIG-IP system to monitor a second mailbox. Continue with #2.

Yes

Select this option if you want the BIG-IP system to monitor a second mailbox.

a. What email address do you want to use for the advanced monitors?

This question only appears if you specified you are deploying Autodiscover and/or Exchange Web Services .

Type the email address associated with the account you are going to monitor (and that you specify in the following question).

b. Which mailbox account should be used for the second monitor?

Type a mailbox account for use in the second monitors. Again, this name corresponds to the account name field in Active Directory (rather than the email address).

c. What is the password for that mailbox account?

Type the associated password.

d. What is the domain name of the user account for the second monitors?

Type the Domain name for the second user account. This domain can be entered in either FQDN (mydomain. example.com) or NetBIOS (MYDOMAIN) format.

2. Which authentication method have you configured for OWA?

This question only appears if you specified you are deploying Outlook Web App.

If you configured the iApp to deploy Outlook Web App at this time, choose the authentication method you have configured for Outlook Web App. The health monitors will be customized to accommodate the authentication method you are using.

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

OWA uses the default Forms-Based authentication

Select this option if you are using Forms-based authentication, which is the default authentication mechanism for OWA.

If you chose Forms-based authentication, the BIG-IP system does not perform an actual login to the service, but checks the availability of the forms-based authentication page.

• OWA uses Basic or Windows Integrated authentication

Select this option if you are using Basic/Windows Integrated authentication.

3. How many seconds should pass between health checks?

Specifies how often the system checks the health of the servers. We recommend the default of 30 seconds. The maximum value for the interval is 28.799 seconds.

4. What FQDN do you use for the OWA service?

Specify the fully qualified domain name you use for your Outlook Web App service.

5. What FQDN do you use for the Outlook Anywhere service?

Specify the fully qualified domain name you use for your Outlook Anywhere service.

6. What FQDN do you use for the ActiveSync service?

Specify the fully qualified domain name you use for your ActiveSync service.

7. What FQDN do you use for the Autodiscover service?

Specify the fully qualified domain name you use for your Autodiscover service.

Additional Steps

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

Finished

Review your answers to the questions. When you are satisfied, click the **Finished** button. The BIG-IP system creates the relevant objects. Continue with *Next steps on page 57*.

Configuring the BIG-IP APM to provide secure remote access to Client Access Servers

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

As mentioned in the prerequisites, because you are deploying BIG-IP APM, you must have configured the BIG-IP system for DNS and NTP. See *Configuring DNS and NTP settings on page 65* for instructions.

Analytics

This section of the template asks questions about Analytics. The Application Visibility Reporting (AVR) module allows you to view statistics specific to your Microsoft Exchange implementation. AVR is available on all BIG-IP systems v11 and later, however you must have the AVR provisioned for this option to appear. Note that this is only for application visibility reporting, you can view object-level statistics from the BIG-IP without provisioning AVR.

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

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

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

Choose whether you want to enable AVR for Analytics.

· No, do not enable Analytics

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

• Yes, enable Analytics using AVR

If you choose to enable Analytics, select Yes from the list, and then answer the following questions.

a. Use the default Analytics profile or select a custom profile?

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

Select a custom Analytics profile

Select this option if you have already created a custom Analytics profile for Exchange Server.

a. Which Analytics profile do you want to use?
 From the list, select the appropriate Analytics profile.

· Use default profile

Select this option if you have not yet created a custom Analytics profile for Microsoft Exchange. We do not recommend using the default profile.

BIG-IP Access Policy Manager

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

1. Would you like to create a new Access Profile, or use an existing Access Profile?

Choose whether you want the system to create a new BIG-IP APM Access Profile, or if you have already created a custom Access Profile outside the template. If you are unsure, select **Create a new Access Profile**.

Select the Access profile you created from the list

If you have previously created an Access profile for your Exchange implementation, select the existing profile you created from the list. Continue with the next section.

· Create a new Access profile

Select this option if you have not created a custom Access profile, and want the system to create one.

a. Would you like to create a new AAA server, or use an existing AAA server?

Choose whether you want the system to create a new BIG-IP APM AAA Server object, or if you have already created a custom AAA Server outside the template. The AAA server contains information about your Active Directory implementation. If you are unsure, select **Create a new AAA Server**.

· Select the AAA Server you created from the list

If you have previously created an AAA Server for your Exchange implementation, select the existing object you created from the list.

a. What is the FQDN of your Active Directory domain for your Exchange users?

Specify the FQDN of the Active Directory deployment for your Exchange users. This is the FQDN for your entire domain, such as example.com, rather than the FQDN for any specific host. Continue with b on the following page.

Create a new AAA Server

Select this option if you have not created a custom AAA Server, and want the system to create one.

a. What is the FQDN of your Active Directory domain for your Exchange users?

Specify the FQDN of the Active Directory deployment for your Exchange users. This is the FQDN for your entire domain, such as example.com, rather than the FQDN for any specific host.

b. Which Active Directory servers in your domain can this BIG-IP system contact?

Specify both the FQDN and IP address of each Active Directory server you want the BIG-IP APM to use for servicing authentication requests. Click **Add** to include additional servers.

c. Does your Active Directory domain allow anonymous binding?

Select whether anonymous binding is allowed in your Active Directory environment.

• Yes, anonymous binding is allowed

Select this option if anonymous binding is allowed. No further information is required.

· No, credentials are required for binding

If credentials are required for binding, you must specify an Active Directory user name and password for use in the AAA Server.

a. Which Active Directory user with administrative permissions do you want to use?

Type a user name with administrative permissions.

b. What is the password associated with that account?

Type the associated password.

d. How do you want to handle health monitoring for this pool?

Specify whether you want the template to create a new LDAP monitor or a new ICMP monitor, or if you select an existing monitor. For more accurate monitoring, we recommend using an LDAP monitor.

Select an existing monitor for the Active Directory pool

Select this option if you have already created a health monitor (only monitors with a Type of LDAP or External can be used) for the Active Directory pool that will be created by the template. If you want to create a health monitor, but have not already done so, you must exit the template and create the object before it becomes available from the list.

The iApp allows you to select monitors that are a part of another iApp Application Service. If you select a monitor that is a part of another Application Service, be aware that any changes you make to the monitor in the other Application Service will apply to this Application Service as well.

a. Which monitor do you want to use?

From the list, select the LDAP or External monitor you created to perform health checks for the Active Directory pool created by the template. Only monitors that have a Type value of LDAP or External appear in this list. Continue with the next question.

Use a simple ICMP monitor for the Active Directory pool

Select this option if you only want a simple ICMP monitor for the Active Directory pool. This monitor sends a ping to the servers and marks the server UP if the ping is successful. Continue with b.

Create a new LDAP monitor for the Active Directory pool

Select this option if you want the template to create a new LDAP monitor for the Active Directory pool. You must answer the following questions:

a. Which Active Directory user name should the monitor use?

Specify an Active Directory user name for the monitor to use when attempting to log on as a part of the health check. This should be a user account created specifically for this health monitor, and <u>must</u> be set to never expire.

b. What is the associated password?

Specify the password associated with the Active Directory user name.

c. What is the LDAP tree for this user account?

Specify the LDAP tree for the user account. As noted in the inline help, ADSI editor, an tool for Active Directory LDAP administration, is useful for determining the correct LDAP tree value. For example, if the user name is 'user1' which is in the organizational unit 'Exchange Users' and is in the domain 'exchange. example.com', the LDAP tree would be: ou=Exchange Users, dc=Exchange, dc=example, dc=com.

d. <u>Does your Active Directory domain require a secure protocol for communication?</u>

Specify whether your Active Directory implementation requires SSL or TLS for communication, or does not require a secure protocol. This determines the port the health monitor uses.

No, a secure protocol is not required

Select this option if your Active Directory domain does not require a secure protocol.

· Yes, SSL communication is required

Select this option if your Active Directory domain requires SSL communication. The health check uses port 636 as the Alias Service Port.

• Yes, TLS communication is required

Select this option if your Active Directory domain requires TLS communication. The health check uses port 389 as the Alias Service Port.

e. How many seconds between Active Directory health checks?

Specify how many seconds the system should use as the health check Interval for the Active Directory servers. We recommend the default of 10 seconds.

b. What text should appear in the user access logon prompt?

Type the text you want users to see above the user name and password prompts when logging on to the BIG-IP APM. By default, this includes the HTML

tag to insert a line break between 'Secure Logon' and 'for F5 Networks'.

Tell us about your Access Policy Manager deployment

This section of the iApp asks about your BIG-IP Access Policy Manager deployment.

1. What IP address do you want to use for the BIG-IP APM virtual server?

Specify the IP address you want to use for the BIG-IP Access Policy Manager virtual server. This is the address clients will use to access the HTTP-based Client Access services.

2. Do you want to re-encrypt the traffic that will be forwarded to your BIG-IP LTM?

Select whether you want the system to re-encrypt traffic that will be sent from this BIG-IP APM to the BIG-IP LTM.

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

• Re-encrypt (SSL Bridging)

Select this option if your implementation requires encrypted traffic to the Client Access Servers, or you are using Exchange 2013 and do not have SP1 or later. The system unencrypts, then re-encrypts the traffic headed for the Client Access Servers.

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

The iApp can create a new Client SSL profile, or if you have created a Client SSL profile which contains the appropriate SSL certificate and key for your Exchange implementation, you can select it from the list.

• Select the Client SSL profile you created from the list

If you manually created a Client SSL profile, select it from the list, and then continue with #2.

Create a new Client SSL profile

Select this option if you want the iApp to create a new Client SSL profile.

a. Which SSL certificate do you want to use?

Select the SSL certificate you imported onto the BIG-IP system for decrypting client connections. If you have not yet imported a certificate, you can leave the default selections and reconfigure this iApp after obtaining the certificates. The deployment will not function correctly until you have selected the correct certificates.

→ Λ

Note: Any certificate that you obtain with multiple names must be in SAN (Subject Alternative Name) format, not SNI (Server Name Indication) format. For more information on SAN certificates, see Subject Alternative Name (SAN) SSL Certificates on page 113.

f. Which SSL key do you want to use? Select the associated key from the list.

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

Select whether you want the iApp to create an F5 recommended Server SSL profile, or if you want to choose a Server SSL profile you already created.

· Select the Server SSL profile you created from the list

If you have previously created a Server SSL profile for your Exchange implementation, from the list, select the existing Server SSL profile you created.

· Create a new Server SSL profile

Select this option if you want the iApp to create a new Server SSL profile.

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

c. Is the remote BIG-IP LTM receiving this traffic using a self-signed or default certificate for decryption, or is the certificate signed by a CA?

Select whether the remote BIG-IP LTM receiving the traffic is using a self-signed (or default) certificate for decrypting the traffic from this system, or if the certificate is signed by a Certificate Authority. Your answer determines the Secure Renegotiation setting on the Server SSL profile. This BIG-IP system will not trust the remote BIG-IP default or a self-signed certificate unless specifically configured to do so in this question.

(i) Important This question pertains to the certificate used by the remote BIG-IP LTM, NOT the certificates present and assigned on the local BIG-IP system you are configuring.

· Certificate Authority-provided certificate and key

Select this option if the remote BIG-IP LTM is using a certificate from a Certificate Authority.

Self-signed or default certificate and key

Select this option if the remote BIG-IP LTM is using a self-signed or default certificate.

Do not re-encrypt (SSL Offload)

Select this option if you do not want the system to re-encrypt traffic to the BIG-IP LTM virtual server. We recommend not re-encrypting unless you have a requirement for SSL for the entire transaction. In this case, the system is offloading the BIG-IP LTM from also having to process the SSL transaction.

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

The iApp can create a new Client SSL profile, or if you have created a Client SSL profile which contains the appropriate SSL certificate and key for your Exchange implementation, you can select it from the list.

Select the Client SSL profile you created from the list

If you manually created a Client SSL profile, select it from the list, and then continue with #2.

· Create a new Client SSL profile

Select this option if you want the iApp to create a new Client SSL profile.

a. Which SSL certificate do you want to use?

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

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

Note: Any certificate that you obtain with multiple names must be in SAN (Subject Alternative Name) format, not SNI (Server Name Indication) format. For more information on SAN certificates, see Subject Alternative Name (SAN) SSL Certificates on page 113.

b. Which SSL key do you want to use?

Select the associated key from the list.

3. What is the virtual IP address on the remote BIG-IP system to which you will forward traffic?

Type the IP address of the virtual server on the remote BIG-IP LTM to which you will be forwarding Client Access traffic from this BIG-IP device. This BIG-IP APM sends traffic to this address after performing authentication.

4. How should the system optimize client-side TCP connections to the BIG-IP LTM?

Select how the system should optimize client-side TCP connections. The iApp uses your selection to configure the proper TCP optimization settings on the TCP profile.

Optimize TCP connections for WAN clients

Select this option if most Exchange server clients are coming into your Exchange environment over a Wide Area Network.

Optimize TCP connections for WAN clients

Select this option if most Exchange server clients are coming into your Exchange environment over a Local Area Network.

a. Which type of authentication do Outlook Web App clients use?

This question only appears if you selected to use BIG-IP APM and for the iApp to create a new Access profile

Choose whether your outlook Web App clients are using Forms-based authentication or Smart Card authentication. You must be using BIG-IP APM version 11.3 or later for Smart Card authentication support for OWA.

· Outlook Web App not used or clients use Forms-based auth

Select this option if your Outlook Web App clients are using Forms-based authentication, or if you are not using OWA.

- a. Would you like to display the OWA computer type and light version options on the APM logon page? Choose whether you want to display the computer type (public/shared vs private) and light version (Use the light version of Outlook Web App) options for OWA on the APM logon page.
 - No, do not display the OWA logon options
 Select this option if you do not want to display the OWA logon options on the APM logon page.
 - · Yes, display the OWA logon options

Select this option if you want users to see the OWA logon options on the BIG-IP APM logon page.

Note that in Exchange 2013, you must enable the logon page options by running a specific PowerShell command in the Exchange Management Shell prior to logging into OWA. See <u>Powershell command for enabling the OWA logon options on page 87.</u>

• Outlook Web App clients use Smart Card authentication

Select this option if your OWA clients use Smart Card authentication and you are using BIG-IP APM v11.3 or later.

- a. Specify a certificate from a CA trusted by this BIG-IP system for client-side processing of smart card authentication. Select the certificate you imported onto the BIG-IP system that is from a Certificate Authority and is trusted by the BIG-IP system for client-side processing of smart card authentication. This certificate must already be imported onto the system before you can select it.
- 5. Should BIG-IP APM restrict EAC access to members of the Exchange Organization Management Security Group?

 Exchange 2013 only This question only appears if you selected Exchange 2013 as your version of Exchange and selected to provide secure authentication with BIG-IP APM.

Select whether you want the BIG-IP APM to restrict Exchange Administration Center (EAC) access to members of Exchange 2013's Organizational Management group. The BIG-IP APM module queries Active Directory group membership for the user making the request to EAC. If the user is not a member of the Organization Management group, the APM policy denies access.

No, do not restrict EAC access by group membership

Select this option and the BIG-IP APM will not restrict access to the EAC by group membership.

• Yes, restrict EAC access by group membership

Select this option if you want to restrict EAC access to the Organization Management group. This adds an additional layer of security to your Exchange deployment, as the system denies access to the EAC from anyone who is not a member of the Organization Management group.

6. Which type of authentication do Outlook Anywhere clients use?

Choose whether your Outlook Anywhere clients use Basic or NTLM authentication. Beginning in BIG-IP version 11.3, the iApp supports using NTLM authentication for Outlook Anywhere.

• Outlook Anywhere clients use Basic Authentication

Select this option if your Outlook Anywhere clients use Basic Authentication. Continue with #5.

Outlook Anywhere clients use NTLM authentication

Select this option if your Outlook Anywhere clients use NTLM information. You must answer the following questions about your Active Directory implementation. Also see <u>Appendix E: Active Directory and Exchange Server configuration for NTLM on page</u> 115 for important information and modifications for NTLM.

i Important Before completing this section, you must create a user account in the same domain that has been properly configured for NTLM delegation. You must create an NTLM Machine Account object on the BIG-IP system to join this system to the Active Directory domain. See <u>Creating an NTLM Machine Account on page 66</u>.

a. Which NTLM machine account should be used for Kerberos delegation?

Select the NTLM Machine Account you created to join the BIG-IP system to the Active Directory domain. If you have not already created an NTLM Machine Account on the BIG-IP system, see <u>Creating an NTLM Machine Account on page</u> <u>66</u>. You must either exit the template now and start over once you have created the NTLM Machine Account, or choose Outlook Anywhere Clients use Basic Authentication from the previous question, and then re-enter the template later.

b. What is the Kerberos Key Distribution Center IP or FQDN?

Specify the IP address or fully qualified domain name of the Kerberos Key Distribution Center (KDC). If you type an FQDN, the BIG-IP system must be able to resolve the address. Otherwise, use the IP address.

c. What is the name of the Kerberos Realm?

Specify the name of the Kerberos Realm. While this name should be in all capital letters, the iApp automatically turns any lower case letters to capital.

d. What is the user name for the Active Directory delegation account you created?

Specify the user name for the Active Directory delegation account you created. This account must be correctly configured in Active Directory for Kerberos delegation. See <u>Appendix E: Active Directory and Exchange Server configuration for NTLM on page 115</u> details.

e. What is the associated password?

Specify the password associated with the account.

7. Do you want to add any iRules to this configuration?

You have the option of adding existing iRules to the virtual server. iRules allow an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. For more information on iRules, see https://devcentral.f5.com/HotTopics/iRules/tabid/1082202/Default.aspx.

i Important Improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your system. Verify the impact of an iRule prior to deployment in production.

If you want to add iRules, from the Options box, select the iRule(s) you want to include, and then click the Add (<<) button.

Additional Steps

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

Finished

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

Modifying the iApp configuration

This section contains modifications you must make to the configuration after running the iApp. Not all of these changes are required in all cases; make sure the change applies to your configuration before modifying the configuration.

Adding iRules to the configuration if you chose to use different IP address for the different CAS services

If you configured the iApp template to use different IP address for the different Client Access Services, and are using ActiveSync and/ or Outlook Anywhere, you must add an iRule to the virtual server(s).

Creating the ActiveSync iRule

If you deployed the iApp for separate virtual servers and are deploying ActiveSync, create the following iRule.

To create the iRule, on the Main tab click iRules > Create. Give the iRule a unique name, and then in the **Definition** field, copy and paste the following code.

```
when HTTP_REQUEST {
COMPRESS::disable
CACHE::disable
}
```

Creating the Outlook Anywhere iRule

If you deployed the iApp for separate virtual servers and are deploying Outlook Anywhere, create the following iRule. To create the iRule, on the Main tab click **iRules > Create**. Give the iRule a unique name, and then in the **Definition** field, copy and paste the following code.

```
1
    when HTTP_REQUEST {
2
        COMPRESS::disable
        CACHE::disable
3
4
5
    when HTTP RESPONSE {
6
7
        if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate"} {
            ONECONNECT::reuse disable
8
            ONECONNECT::detach disable
9
            NTLM::disable
10
11
        }
12
            if {[HTTP::header exists "Transfer-Encoding"]} {
13
            HTTP::payload rechunk
        }
14
15
   }
```

To attach the iRule(s) to the virtual server

- 1. From the Main tab of the BIG-IP Configuration utility, expand iApp and then click Application Services.
- 2. Click the name of your existing Microsoft Exchange application service from the list.
- 3. On the Menu bar, click Reconfigure.
- 4. If necessary, from the Do you want to customize your server pool settings? question, select Customize pool settings.
- 5. If you created the ActiveSync iRule, from the *Do you want to add any iRules to this virtual server?* question under the question asking if you are deploying ActiveSync, select the iRule you just created and then click the Add (<<) button to move it to the **Selected** list.
- 6. If you created the Outlook Anywhere iRule, from the *Do you want to add any iRules to this virtual server?* question under the question asking if you are deploying Outlook Anywhere, select the iRule you just created and then click the Add (<<) button to move it to the **Selected** list.
- 7. Click Finished.

Optional: Configuring the BIG-IP system to support MAPI over HTTP in Exchange 2013 SP1

Introduced in Exchange 2013 SP1, the new MAPI over HTTP transport protocol is for Outlook clients running Office 2013 SP1 and later (only). This new service is not yet included in the iApp template, so you must manually configure the BIG-IP system to support it.

If you are using Microsoft Exchange 2013 SP1 or later and using the new MAPI over HTTP transport protocol, use the following guidance to create the objects necessary to support MAPI over HTTP. If you configured the iApp template to use a combined virtual server, you create a health monitor, pool, and an iRule.



(i) Important Because BIG-IP APM is not yet supported for MAPI over HTTP, the iRule in the following table includes a line (commented out by default) to disable Access Policy processing for this new protocol only. If you configured the iApp to use separate virtual servers, you create the monitor, pool, and a virtual server. The iRule is not necessary at all in this case.

Use the following table to create the objects on the BIG-IP LTM. Unless otherwise specified, settings not mentioned in the table can be configured as applicable for your configuration. For instructions on configuring individual objects, see the online help or product manuals.

Health Monitors (Main ta	ab>Local Traffic>Monito	ors)		
Simple health monitor				
Name	Type a unique name			
Туре	HTTP (if using SSL offload) or HTTPS (if using SSL bridging)			
Interval	30 (recommended)			
Timeout	91 (recommended)			
Send String	GET /mapi/healthcheck.htm HTTP/1.1\r\nHost: mapi.example.local\r\nConnection: Close\r\n\r\n			
Receive String	200 OK	200 OK		
Advanced Monitor				
The advanced monitor for MAPI next section and set the Availab		onitor as the advanced monitor for the EWS service. You simply add the EWS monitor to the pool in the cribed.		
Pools (Main tab>Local Tr	affic>Pools)			
Name	Type a unique name			
Health Monitor	Select the monitor you created above. If you are using the advanced monitor, add both the advanced and simple monitor.			
Availability Requirement	If using the advanced monitor (only), select All			
Load Balancing Method	Least Connections (Member)			
Address	Type the IP Address of your server			
Service Port	80 (if using SSL offload)	or 443 (if using SSL bridging) Click Add to repeat Address and Service Port for all nodes		
Profiles (Main tab>Local	Traffic>Profiles)			
UTTO	Parent Profile	http		
HTTP	Redirect Rewrite	Matching		
TCP WAN¹	Parent Profile	tcp-wan-optimized		
TCP LAN¹	Parent Profile	tcp-lan-optimized		
	Parent Profile	clientssl		
Client SSL	Certificate/Key	Select the Certificate and Key you imported		
	Parent Profile	serverssl		
Server SSL ²	Options List	If using BIG-IP v11.4.x only, enable No TLSv1.2		
	Parent Profile	oneconnect		
OneConnect	Source Mask	255.255.255.255		
NTLM	Parent Profile	ntim		
iRules (Main tab>Local T	raffic>iRules) This iRule	e is for the <u>combined</u> virtual server scenario only		
Name	Type a unique name			
Definition	See the following section	for the iRule definition and instructions on attaching the iRule to the virtual server using the iApp.		

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.

Virtual Servers (Main tab-->Local Traffic -->Virtual Servers) This virtual server is only for the separate virtual server scenario Destination Address IP address for the virtual server Service Port 443 Profiles Add each of the profiles you created from the appropriate list Secure Address Translation iRules If using a combined virtual server, Default Pool Select the pool you created for MAPI over HTTP

Creating the iRule definition for the combined virtual server scenario

Use the following for the Definition of the iRule, omitting the line numbers, and *changing the red text to the name your pool*. If you want MAPI over HTTP to bypass the BIG-IP APM, remove the comment (#) from line 5.

```
when HTTP_REQUEST {
2
       switch -glob -- [string tolower [HTTP::path]] {
3
          "/mapi*" {
             ###uncomment the following line to bypass APM for MAPI-over-HTTP
4
5
             #ACCESS::disable
6
             pool mapi_http_pool
             COMPRESS::disable
7
8
             CACHE::disable
9
             return
10
          }
11
       }
12
    }
```

To attach the iRule to the combined virtual server

- 1. From the Main tab of the BIG-IP Configuration utility, expand iApp and then click Application Services.
- 2. Click the name of your existing Microsoft Exchange application service from the list.
- 3. On the Menu bar, click Reconfigure.
- 4. If necessary, from the Do you want to customize your server pool settings? question, select Customize pool settings.
- 5. From the Do you want to add any iRules to this combined virtual server? question, select the iRule you just created and then click the Add (<<) button to move it to the **Selected** list.
- 6. Click Finished.

Optional: Configuring APM to Support Windows Integrated Authentication For Outlook Web App

By default, the Exchange iApp template supports deploying APM for pre-authentication to Outlook Web App when OWA is configured for Forms-based authentication.

If you are deploying BIG-IP APM and have configured Outlook Web App for Windows Integrated authentication, use the following guidance for configuring the iApp template, and then modifying the configuration.

- 1. Configure the iApp template as applicable for your implementation with the following exceptions:
 - a. For the question Which type of authentication do Outlook Web App clients use? select Outlook Web App clients use Forms-based authentication.
 - b. If using advanced health monitors, for the question *Which authentication method have you configured for OWA?* select **OWA uses Basic or Windows Integrated authentication**.
- 2. Disable the Strict Updates feature:
 - a. Click iApps > Application Services > name you gave the Exchange application service.
 - b. On the Menu bar, click Properties.
 - c. In the **Strict Updates** field, clear the box to disable Strict Updates. You may have to select **Advanced** from the **Application Service** list at the top of the box to see this option.
 - d. Click Update.
- 3. Modify the select SSO iRule:
 - a. Click iRules > (name you gave the Exchange application service)_select_sso_irule.
 - b. In the **Definition** section, copy and paste the following, replacing **<app_name>** in line 4 with the name you gave the iApp.

```
when ACCESS_ACL_ALLOWED {
    set req_uri [string tolower [HTTP::uri]]
    if { $req_uri contains "/owa" } {
        WEBSSO::select [set foo /Common/<app_name>.app/exch_ntlm_sso]
    }
    unset req_uri
}
```

- c. Click Update.
- 4. You must edit the existing OWA timeout iRule to automatically log users out. See <u>Optional: Creating the iRule to terminate</u> <u>inactive APM sessions if using Windows based authentication for OWA on page 95</u>. Note this is **not** optional if you used the iApp to configure your Exchange implementation.

Optional: Configuring BIG-IP LTM/APM to support NTLMv2-only deployments

If you have configured your Microsoft Windows domain to support only NTLMv2 authentication and refuse LM/NTLM requests, you must either modify the configuration produced by the template by disabling the Strict Updates feature, or create a new APM profile manually and then assign the profile to the configuration using the iApp template.

Choose one of the following procedures.

Manually creating an APM profile

Use the BIG-IP APM manual configuration table on pages <u>BIG-IP APM Configuration on page 87</u> to create the APM objects. Where applicable, use the NTLMv2 option.

Once you have created the APM Access Profile and associated objects, you can reconfigure the iApp and select the Access Profile you created.

- 1. Re-enter the iApp template (on the Main tab, click iApp > Application Services > [name of your Exchange application service] and then from the Menu bar, click Reconfigure).
- 2. In the BIG-IP Access Policy Manager (APM) section, from the "Would you like to create a new Access Profile, or use an existing Access Profile?" list, select the profile you just created.
- 3. Click Update.

This completes the modifications for NTLMv2 if you manually configured the APM profile.

Disabling strictness on the iApp deployment

If you do not want to create an entire new APM profile with the associated objects, after deploying the template, you can disable the Strict Updates feature on the iApp and modify existing objects to support NTLMv2. You will need to create an NTLMv2 SSO object manually, and then modify the Exchange APM Profile produced by the template to reference that SSO configuration.

- 1. Use the BIG-IP APM manual configuration table on pages <u>BIG-IP APM Configuration on page 87</u> to create the NTLMv2 SSO Configuration object.
- 2. Disable the Strict Updates feature on the iApp Application Service (see step 2 of the procedure page 55).
- 3. The next step depends on which version of the BIG-IP system you are using:
 - BIG-IP v11.3 or earlier
 On the Main tab, click Access Policy > Access Profile > Name of the Access Profile created by the template From the SSO Configuration list, select the NTLMv2 object you created.
 - BIG-IP v11.4 or later
 On the Main tab, click Access Policy > Application Access > Microsoft Exchange > Name of the Access Profile created by the template > Edit (BIG-IP v11.4 and later). Under Service Settings on the left, click each of the Exchange Services, and from the SSO Configuration list, where an NTLM SSL Configuration object is selected, select the NTLMv2 object you created.
- **i** Important F5's external monitors for Autodiscover, Outlook Anywhere, and EWS do not support NTLMv2. If you have configured your domain to refuse LM and NTLM requests, you must select "Use simple monitors" in response to the "Do you want to use advanced or simple server health monitors?" question in the Server Health Monitors section of the template.

Next steps

After completing the iApp Template, the BIG-IP Application Services page opens for the Exchange application service you just created. To see the list of all the configuration objects created to support Microsoft Exchange, on the Menu bar, click **Components**. The complete list of all Exchange related objects opens. You can click individual objects to see the settings.

Once the objects have been created, you are ready to use the new deployment.

Modifying DNS settings to use the BIG-IP virtual server address

Before sending traffic to the BIG-IP system, your DNS administrator may need to modify any DNS entries for the Exchange implementation to point to the BIG-IP system's virtual server address.

Modifying the iApp configuration

The iApp application service you just created can be quickly and easily modified if you find it necessary to make changes to the configuration. The Strict Updates feature of the iApp prevents users from manually modifying the iApp configuration (Strict Updates can be disabled, but use extreme caution). iApp allows you to re-enter the template, make changes, and then update the template. The modifications are automatically made to any of the associated objects.

To modify the configuration

- 1. On the Main tab, expand iApp and then click Application Services.
- 2. Click the name of your Exchange Application service from the list.
- 3. On the Menu bar, click Reconfigure.
- 4. Make the necessary modifications to the template.
- 5. Click the Finished button.

Viewing statistics

You can easily view a number of different statistics on the BIG-IP system related to the Exchange configuration objects created by the iApp template. You can get statistics specific to the Application Service if you have provisioned AVR. Otherwise, you can always get object-level statistics.

AVR statistics

If you have provisioned AVR, you can get application-level statistics for your Exchange application service.

To view AVR statistics

- 1. On the Main tab, expand iApp and then click Application Services.
- 2. From the **Application Service** List, click the Exchange 2010 service you just created.
- 3. On the Menu bar, click Analytics.
- 4. Use the tabs and the Menu bar to view different statistics for your Exchange iApp.

Object-level statistics

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

To view object-level statics

- 1. On the Main tab, expand **Overview**, and then click **Statistics**.
- 2. From the Statistics Type menu, you can select Virtual Servers to see statistics related to the virtual servers.
- 3. You can also choose **Pools** or **Nodes** to get a closer look at the traffic.
- 4. To see networking statistics in a graphical format, click **Dashboard**.

For more information on viewing statistics on the BIG-IP system, see the online help or product documentation.

Adding Ratio or Connection Limit information to the nodes if using a node-based load balancing method

If you chose to customize the server pool settings, changed the load balancing method from the default to a node-based method (such as Ratio (node) or Least Connections (node)), <u>and</u> configured a Ratio or Connection Limit, the iApp applies the ratio or connection limit to the load balancing pool member, and not to the node itself. In this case, you must manually modify each node to include any Ratio or Connection Limit settings you want to configure.

To modify the nodes to include Ratio or Connection Limit settings

- 1. On the Main tab, expand Local Traffic and then click Nodes.
- 2. From the Node table, click a Client Access Server node you entered in the iApp template.
- 3. In the **Ratio** box, type the appropriate ratio, if applicable.
- 4. In the **Connection Limit** box, type the appropriate connection limit, if applicable.
- 5. Click Update.
- 6. Repeat this procedure for each node that is a part of your Exchange deployment.

Troubleshooting

This section contains common issues and troubleshooting steps.

> Advanced health checks are fail when using Windows Integrated Authentication (NTLM provider)

If you are using Windows Integrated Authentication (NTLM provider) only, the BIG-IP health checks using a valid account may fail, as the BIG-IP system fails to correctly form the authentication request headers.

If you are using Windows Authentication with NTLM and you have disabled Basic authentication for the Exchange service you are monitoring, you must manually delete the \r\n at the end of the Send String, and the <domain>\ information from the User Name field.

(i) Important F5 monitors support NTLMv1 authentication. You must ensure that the LmCompatibilityLevel setting in Group Policy for the domain used by the monitor credential is configured to support NTLMv1.

Modifying the IIS authentication token timeout value

The iApp template configures 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):

http://www.microsoft.com/technet/prodtechnol/WindowsServer2003/Library/IIS/6b2e7fcd-5fad-4ac8-ac0a-dcfbe771e9e1.mspx

 \triangle

Warning Use extreme caution any time you are editing the registry. Contact Microsoft for specific instructions and/or help editing the registry values.

> Microsoft Exchange Remote Connectivity Analyzer fails to successfully run the FolderSync command

If you deployed the BIG-IP system for ActiveSync, either using the iApp template or manually, and attempt to run the Microsoft Exchange Remote Connectivity Analyzer (ExRCA) against an Exchange mailbox, you may receive the following error:

Attempting the FolderSync command on the Exchange ActiveSync session.

The test of the FolderSync command failed.
Additional Details: Exception details:
Message: The request was aborted: The request was canceled.
Type: System.Net.WebException
Stack trace:
at System.Net.HttpWebRequest.GetResponse()

at Microsoft.Exchange.Tools.ExRca.Extensions.RcaHttpRequest.GetResponse()

This behavior affects versions of BIG-IP earlier than 11.4.0. To work around this error, you must create an iRule, and then use the iApp template to apply the iRule to the combined Exchange BIG-IP virtual server (or attach the iRule manually if you used the manual configuration tables).

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.
- 3. In the **Definition** section, copy and paste one of the following iRules, omitting the line numbers, depending on whether you configured the system for a combined virtual server, or a separate virtual server for ActiveSync.

Only use the definition applicable to your configuration.

Combined virtual server iRule definition

```
when HTTP_REQUEST {
        set isactivesync 0
2
3
        if { [string tolower [HTTP::path]] contains "/microsoft-server-activesync" } {
4
                set isactivesync 1
5
6
   }
    when HTTP RESPONSE {
7
8
        if { [HTTP::status] == 401 && [HTTP::header exists "Content-Length"] && $isactivesync == 1 } {
        HTTP::header insert "Connection" "Close"
9
10
11
        unset isactivesync
12 }
```

Separate virtual server iRule definition

```
when HTTP_RESPONSE {
    if { [HTTP::status] == 401 && [HTTP::header exists "Content-Length"] } {
        HTTP::header insert "Connection" "Close"
    }
}
```

4. Click Finished.

The next task is to attach the iRule to the virtual server. This depends on whether you configured the BIG-IP system using the iApp template or manually.

Attaching the iRule if you used the iApp template to configure the BIG-IP system

Use the following procedure if you used the iApp template to configure the BIG-IP system.

To attach the iRule to the virtual server

- 1. From the Main tab of the BIG-IP Configuration utility, expand iApp and then click Application Services.
- 2. Click the name of your existing Microsoft Exchange application service from the list.
- 3. On the Menu bar, click **Reconfigure**.
- 4. If necessary, from the *Do you want to customize your server pool settings?* question, select **Customize pool settings**.
- 5. If you used a <u>Combined</u> virtual server, from the *Do you want to add any iRules to this combined virtual server?* question, select the iRule you just created and then click the Add (<<) button to move it to the **Selected** list.

If you used <u>Separate</u> virtual servers, after the question *What IP address do you want to use for the ActiveSync virtual server*? from the *Do you want to add any custom iRules to this virtual server*? question, select the iRule you just created and then click the Add (<<) button to move it to the **Selected** list.

6. Click Finished.

Attaching the iRule if you manually configured the BIG-IP system

If you configured the BIG-IP system manually, and configured a combined virtual server, modify the combined virtual server you created to attach the combined iRule.

If you configured separate virtual servers, modify the ActiveSync virtual server you created to attach the separate virtual server iRule.

Advanced monitors for Autodiscover, EWS, and Outlook Anywhere only support Basic and NTLMv1 authentication

The advanced monitors for Autodiscover, Exchange Web Services, and Outlook Anywhere support Basic and NTLMv1 authentication only. If you have configured your domain to use NTLMv2 only, you must modify the health monitors to remove the **--ntlm** option from the curl statement used in the Autodiscover, EWS, and Outlook Anywhere external monitors (if you deployed the template for these services) using the guidance in this section.

Additionally, you will need to enable Basic authentication for the EWS virtual directory using the IIS Manager snap-in on each Exchange Client Access Server. Consult the Microsoft documentation for instructions.

i Important This is only necessary if you have configured your domain to use NTLMv2

To modify the health monitors

- a. On the Main tab, click System > File Management > External Monitor Program File List.
- b. Depending on which services you deployed, click either **autodiscover_eav**, oa_eav, or ews_eav.

 The file name for Autodiscover is always autodiscover_eav and the file name for Outlook Anywhere is always oa_eav.

 If you deployed EWS without Outlook Anywhere, the file name is ews_eav. Unless you have multiple instances of the iApp, you should never have both an oa_eav and ews_eav.
 - **Note**: If you have multiple instances of the iApp template, make sure you click the file name in the applicable Partition/Path on the far right of the table. The file names are always the same.
- c. Locate the line that begins with **curl** and remove the **--ntlm** portion only. When you are finished, this line should look similar to the following:

```
curl -g -s -k -X POST -H 'Content-Type: text/xml; charset=utf-8' -d "${XMLFULL}" -u
${DOMAIN}\\${USER}:${PASSWORD} https://${NODE}${ADSURI} | grep -i "${RECV}" 2>&1 > /dev/null
```

- d. Click Update.
- e. If you deployed other applicable services, repeat steps b e for any to remove --ntlm from that file.
 - **i** Important If you re-enter the iApp template and modify the configuration using the Reconfigure option, you must make these changes again, as the iApp will overwrite the modifications.

Guest accounts on the BIG-IP system can view the persistence table

This issue was fixed in version 1.3 of the iApp template. You should not experience this issue in current versions.

Because the Exchange iApp uses the Basic authorization header for ActiveSync and Outlook Anywhere session persistence, BIG-IP guest accounts that have been explicitly granted access to the Traffic Management Shell (tmsh) are able to view encoded user credential and password information. It is possible an attacker logging in to BIG-IP as a guest could decode these credentials. F5 recommends disabling tmsh access for any BIG-IP guest accounts by clicking **System > Users > User List > Terminal Access > Disabled > Finished**.

Alternately, you may edit the iRule(s) created by the iApp template to obfuscate the encoded credentials. These changes are not necessary if you have used the iApp template to deploy F5's Access Policy Manager module. Replace all instances of the following:

persist uie [HTTP::header "Authorization"] 7200 with:

set <service>_key [sha256 [HTTP::header "Authorization"]]

persist uie \$<service>_key 7200

Where **<service>** indicates either ActiveSync or Outlook Anywhere. For example:

set oa_key [sha256 [HTTP::header "Authorization"]]

persist uie \$oa_key 7200

iPhones and other iOS devices are displaying invalid certificate messages after deploying the iApp for 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 manually create an Client SSL profile that uses 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.

Use the guidance in this solution to create a Client SSL profile that uses an intermediate certificate chain: http://support.f5.com/kb/en-us/solutions/public/13000/300/sol/13302.html.

Be sure Secure Renegotiation is set to Require (the default) on the Client SSL profile.

If you manually configured the system, add the Client SSL profile to your virtual server.

If you used the iApp, use this procedure:

- a. Re-enter the iApp template (on the Main tab, click iApp > Application Services > [name of your Exchange application service] and then from the Menu bar, click Reconfigure).
- b. In the *Tell us about your deployment* section, from the "Do you want to create a new client SSL profile or use and existing one?" question, select the profile you just created that uses the Chain certificate.
- c. Click Update.

When using SSL Bridging and BIG-IP version 11.4.x, pool members may be marked down or you may experience connection resets and TLS errors logged to the Client Access servers

This issue only occurs when using SSL Bridging <u>and</u> BIG-IP versions 11.4.x. Pool members may be marked down when using simple monitors, or you may experience connection resets and TLS errors logged to the Client Access servers because the SSL ciphers used in the Server SSL profile in 11.4.x are not compatible with those in some versions of Microsoft Internet Information Server (IIS).

There are two ways you can resolve this issue:

- 1. Upgrade your BIG-IP system to version 11.5 or later.
- 2. Create a custom Server SSL profile and associate it with the virtual server, either using the iApp template or manually.

To create the Server SSL profile

- a. On the Main tab, click Local Traffic > Profiles > SSL > Server.
- b. Click Create.
- c. In the **Name** box, type a unique name for this profile.
- d. In the Options row, click the Custom box.
- e. From the Available Options list, select No TLSv1.2, and then click the Enable button.
- f. Click the Finished button.
- g. Attach the new Server SSL profile to the virtual server either using the iApp or manually.
 - To attach the profile to the virtual server using the iApp template:
 - i) Re-enter the iApp template (on the Main tab, click iApp > Application Services > [name of your Exchange application service] and then from the Menu bar, click Reconfigure).
 - ii) In the Tell us which services you are deploying section, from the "Which Server SSL profile do you want to use" question, select the Server SSL profile you just created.
 - iii) Click Update.
 - To attach the profile to the virtual server manually:
 - i) Select the Exchange virtual server you created.
 - ii) From the SSL Profile (Server) area, enable the Server SSL profile you just created.
 - iii) Click Update.
 - iv) If you used separate virtual servers for each Exchange service, add the profile to each virtual server.

Lync clients cannot connect or receive authentication prompts when accessing Microsoft Exchange Autodiscover and EWS through F5 APM

When you have deployed APM in front of Microsoft Exchange 2010 or 2013, Microsoft Lync clients may be unable to successfully query the Autodiscover service or download free/busy information from EWS. To work around this issue, you must create an iRule to disable APM for these requests and attach it using the iApp interface.

To create the iRule and add it to the virtual server

- 1. On the Main tab, click Local Traffic > iRules > Create.
- 2. In the **Name** box, type a name.
- 3. In the **Definition** section, copy and paste the following iRule, omitting the line numbers. Enter line 4 as a single line.

```
priority 1
    when HTTP_REQUEST {
3
        set is_disabled 0
4
        if { [string tolower [HTTP::header value "User-Agent"]] contains "microsoft lync" || [string tolower [HTTP::header
    value "User-Agent"]] contains "ms-webservices"
                                                        } {
5
           if { [string tolower [HTTP::path]] starts_with "/autodiscover" } {
6
              set is_disabled 1
7
              set path [HTTP::path]
8
              ACCESS::disable
              HTTP::path disable-$path
10
              pool <your Autodiscover pool>
11
           if { [string tolower [HTTP::path]] starts_with "/ews" } {
   set is_disabled 1
12
13
14
              set path [HTTP::path]
              ACCESS::disable
15
              HTTP::path _disable-$path
pool <your EWS pool>
16
17
18
              COMPRESS::disable
19
              CACHE::disable
20
           }
21
        }
22
23
    when HTTP_REQUEST_RELEASE {
24
         if { !$is_disabled }
                                { return }
25
         HTTP::path $path
26
         unset is_disabled
27
```

- 4. Click the Finished button.
- 5. Re-enter the iApp template (on the Main tab, click iApp > Application Services > [name of your Exchange application service] and then from the Menu bar, click Reconfigure).
- 6. From the iApp interface, select "Customize pool settings" from the "Tell us about which services you are deploying" section, from the "Do you want to add any iRules to this combined virtual server?" question, enable the iRule you created.
- 7. Click Update.

After attempting to deploy APM for Outlook Anywhere with NTLM authentication using the iApp template, you receive a "Script did not successfully complete" error

This issue was fixed in version 1.5.0rc2 of the iApp template. You should not experience this issue in current versions.

If you are deploying APM with NTLM authentication for Outlook Anywhere, and you use specific characters in the Active Directory delegation account password field, you receive an error and the iApp fails to deploy. You may see errors such as:

- » Script did not successfully complete: (no such variable)
- » Script did not successfully complete: (invalid command name)
- » Script did not successfully complete: (unknown property)

If upgrading is not an option, use a password that does not include the characters \$, [,], #, or;

Clients receiving error message when using BIG-IP APM with OWA 2013 and IE10 or Google Chrome

If you are using BIG-IP APM and Outlook Web App 2013, and have clients using Internet Explorer 10 or Google Chrome, clients may receive the following error message from the BIG-IP APM: *Access policy evaluation is already in progress for your current session.* If clients are receiving this error, you must apply the an iRule to the virtual server(s) used for OWA 2013.

To create the iRule and add it to the OWA 2013 virtual server

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

```
when HTTP_REQUEST {
    if { [HTTP::cookie exists "IsClientAppCacheEnabled"] } {
        HTTP::cookie "IsClientAppCacheEnabled" False
    }
}
```

- 5. Click the Finished button.
- 6. Re-enter the iApp template (on the Main tab, click iApp > Application Services > [name of your Exchange application service] and then from the Menu bar, click **Reconfigure**).
- 7. In the Tell us which services you are deploying section, from the "Do you want to customize your server pool settings" question, select **Customize pool settings**. Either in the "Do you want to add any custom iRules to this combined virtual server?" (if you used a single IP address) or in the "Do you want to add any custom iRules to this virtual server?" question under the IP address for OWA question (if you used different IP addresses), enable the iRule you just created.
- 8. Click Update.

If you have Outlook Web App clients connecting to a BIG-IP APM virtual server externally, and the same clients connect to a non-APM virtual server internally, you must apply the iRule to both virtual servers.

If clients are still receiving this error after adding the iRule, you should request they delete Temporary Internet Files (IE10), or go to chrome://appcache-internals and remove the application cache for Outlook Web Access (Chrome).

You may experience deployment errors when the NTLM Machine Account name contains spaces or special characters

If you are using BIG-IP APM, and specified that Outlook Anywhere clients use NTLM authentication, you must specify an NTLM Machine Account in the iApp template that you created manually. If the name of the NTLM Machine Account object contains spaces or special characters, you may experience errors when trying to deploy the template.

If your NTLM Machine Account object name contains a special character or space, the workaround for this issue is to create an NTLM Machine Account name that only contains alphanumeric characters and underscores, with no spaces. Return to *Creating an NTLM Machine Account on page 66*, and create a new machine account.

The Direct File Access setting for public computers is not honored

When you have configured the OWA virtual directory to deny Direct File Access to public computers, and you have deployed BIG-IP APM with OWA logon options enabled, users who have selected **This is a public or shared computer** from the APM logon page are able to download or open OWA file attachments.

To solve this issue, create and then attach the following iRule to the combined virtual server or the separate OWA virtual server (this rule should appear below the _owa_forms_value_irule in the iRule list):

```
when HTTP_REQUEST {
    if { [ACCESS::session data get "session.custom.owa.trusted"] == 0 } {
    if { [HTTP::cookie exists "PrivateComputer"] } {
        HTTP::cookie remove "PrivateComputer"
}
}
```

Appendix A: Configuring additional BIG-IP settings

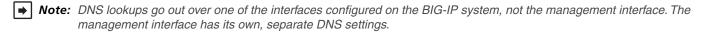
This section contains information on configuring the BIG-IP system for objects or settings that are required, but not part of the template.

Configuring DNS and NTP settings

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

Configuring the DNS settings

In this section, you configure the DNS settings on the BIG-IP system to point to a DNS server that can resolve your Active Directory server or servers. In many cases, this IP address will be that of your Active Directory servers themselves.



(i) Important The BIG-IP system must have a self IP address in the same local subnet and VLAN as the DNS server, or a route to the DNS server if located on a different subnet. The route configuration is found on the Main tab by expanding Network and then clicking Routes. For specific instructions on configuring a route on the BIG-IP system, see the online help or the product documentation.

To configure DNS settings

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

Configuring the NTP settings

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

To configure NTP settings

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

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

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

Creating an NTLM Machine Account

If you are using BIG-IP APM to provide secure authentication and configuring the BIG-IP system for Outlook Anywhere clients using NTLM authentication, you must have an NTLM Machine Account object configured before you can successfully complete the template. Use the following procedure to create the NTLM Machine Account.

To create the NTLM Machine Account

- 1. On the Main tab, expand Access Policy, and then click Access Profiles.
- 2. On the Menu bar, from the NTLM menu, click Machine Account List.
- Click the Create button.
- 4. In the **Name** box, type a name for the BIG-IP Machine Account object. Currently, the NTLM machine account should contain alphanumeric characters and underscores only. Spaces and special characters are not allowed.
- 5. In the **Machine Account Name** box, type the name of the computer account that will be created in the domain after clicking Join.
- 6. In the **Domain FQDN** box, type the fully qualified domain name of the domain that you want the machine account to join.
- 7. In the **Domain Controller FQDN** box, if the machine account should have access to one domain only, type the FQDN for the domain controller for that domain.
- 8. In the **Admin** User box, type the name of a user with administrative privileges.
- 9. In the **Password** box, type the associated password.
- 10. Click the Join button.

Appendix B: Using X-Forwarded-For to log the client IP address

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

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

Deploying the Custom Logging role service

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

To deploy the Custom Logging role service in Windows 2008 and 2008 R2

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

To deploy the Custom Logging role service in Windows 2012 and 2012 R2

- 1. From your Windows Server 2012 or Windows Server 2012 R2 device, open Server Manager.
- 2. Click Add Roles and Features.
- 3. In the Add Roles and Features wizard, the Custom Logging Role is under the **Web Server > Web Server > Health and Diagnostics** category.
- 4. On the Confirmation page, click Install.
- 5. After the service has successfully installed, click the **Close** button.

Adding the X-Forwarded-For log field to IIS

Before beginning the following procedure, you must have installed IIS Advanced Logging. For installation instructions, see http://www.iis.net/community/files/media/advancedlogging_readme.htm

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.
- 6. Click a Log Definition to select it. By default, there is only one: %COMPUTERNAME%-Server. The log definition you select must have a status of Enabled.
- 7. From the Actions pane on the right, click Edit Log Definition.
- 8. Click the Select Fields button, and then check the box for the X-Forwarded-For logging field.
- 9. Click the **OK** button.
- 10. From the Actions pane, click Apply.
- 11. Click Return To Advanced Logging.
- 12. In the Actions pane, click Enable Advanced Logging.

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

Appendix C: Manual configuration tables

This table contains the BIG-IP configuration objects in this deployment and any non-default settings. See the BIG-IP APM tables for additional APM configuration. Give each BIG-IP object a unique name in the Name field. Because of the complexity, we strongly recommend using the iApp to configure Microsoft Exchange Server. For the new MAPI over HTTP service in Exchange 2013 SP1, see Optional: Configuring the BIG-IP system to support MAPI over HTTP in Exchange 2013 SP1 on page 53.

Note: We recommend using this section to create two monitors for each service, using a second mailbox account.

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

Health Monitors (Main tab > Local Traffic > Monitors)				
Outlook Web App monitor (includes ECP)				
Туре	HTTP (SSL offload), HTTPS (SSL Bridging). If using Exchange 2013, you must use HTTPS.			
Interval	30 (recommended)			
Timeout	91 (recommended)			
Send String 1	If using the default forms-based authentication for OWA:			
		<pre>gon.aspx?url=https://mail.example.com/owa/&reason=0 HTTP/1.1\r\nUser-Agent: st: mail.example.com\r\n</pre>		
	If using Basic or Basic	c and WIndows Integrated Authentication for OWA:		
	GET /owa/ HTTP/1.	1\r\nUser-Agent: Mozilla/4.0\r\nHost: mail.example.com\r\n		
Receive String 2	Exchange Server 2010: OutlookSession = Exchange Server 2013: 200 OK			
User Name	Type the appropriate us	ser name of a valid mailbox account.		
Password	Type the associated pa	ssword		
ActiveSync monitor				
Туре	HTTP (SSL offload), HTTPS (SSL Bridging). If using Exchange 2013 SP1, you must use HTTPS.			
Interval	30 (recommended)			
Timeout	91 (recommended)			
Send String 1	OPTIONS /Microsof	t-Server-ActiveSync/ HTTP/1.1\r\nHost: mail.example.com\r\n		
Receive String	MS-ASProtocolCommands: Sync,SendMail,SmartForward,SmartReply,GetAttachment,GetHierarchy,CreateCollection,DeleteCollection,MoveCollection,FolderSync			
User Name	Type the appropriate user name of a valid mailbox account.			
Password	Type the associated pa	ssword		
Outlook Anywhere monitor	r (includes EWS)			
Туре	External			
Interval	30 (recommended)			
Timeout	91 (recommended)	91 (recommended)		
External Program	See Importing the monitor script files on page 79 for the EAV script			
Variables	Name	Value		
	USER PASSWORD DOMAIN EMAIL	The account name associated with a mailbox. The password for the account The Windows domain for the account The email address for the user mailbox (such as j.smith@example.com)		
Autodiscover monitor				
Туре	External			
Interval	30 (recommended)			
Timeout	91 (recommended)			
External Program	See Importing the monitor script files on page 79 for the EAV script			
Variables	Name	Value		
	USER PASSWORD DOMAIN EMAIL	The account name associated with a mailbox. The password for the account The Windows domain for the account The email address for the user mailbox (such as j.smith@example.com)		

¹ For Advanced Monitors only. Simple monitors only require the Type, Interval, and Timeout. Replace red text with your FQDN. It must be on a single line.

² 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.

³ You must only include a single \r\n at the end of the string.

Pools (Main tab>Local Tra	affic>Pools) Important:	Repeat for each Client Access Server role	
Health monitor		onitor 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 (443 if using SSL Bridging	Repeat Address and Port for all members	
iRules (Main tab > Local Tra	affic > iRules)		
iRules	OWA Redirect iRule	Create the Redirect iRule, using the Definition found on page 80	
(Local Traffic>iRules)	Persistence iRule	Create the Persistence iRule, using the Definition found on page 80	
Profiles (Main tab > Local 7	Traffic > Profiles)		
НТТР	Parent Profile	http	
(Profiles->Services)	Redirect Rewrite	All	
HTTP Compression (Profiles->Services)	Content List>Include List	See HTTP Compression Content include list on page 85	
Web Acceleration (Profiles >Services)	Parent Profile	optimized-caching	
TCP WAN¹ (Profiles >Protocol)	Parent Profile	tcp-wan-optimized	
TCP LAN¹ (Profiles->Protocol)	Parent Profile	tcp-lan-optimized	
Client SSL	Parent Profile	clientssl	
(Profiles->SSL)	Certificate/Key	Select the Certificate and Key you imported	
Server SSL ²	Parent Profile	serverssl	
(Profiles->SSL)	Options List	If using BIG-IP v11.4.x only, enable No TLSv1.2	
Persistence (Profiles->Persistence)	Persistence Type	Cookie (Exchange 2010 only) Important: If using 2010, RPC Client Access, and OAB/EWS see Adding the persistence profiles if using Exchange 2010, RPC Client Access, and OAB/EWS on page 78	
OneConnect	Parent Profile	oneconnect	
(Profiles->Other)	Source Mask	255.255.255.255	
NTLM (Profiles->Other)	Parent Profile	ntlm	
Virtual Servers (Main tab >	> Local Traffic > Virtual Serv	ers)	
	Destination Address	IP address for the virtual server (Service Port 443)	
	Profiles	Add each of the profiles you created above from the appropriate list	
Port 443	SNAT Pool ³	Auto Map³	
	iRules	Add the Append and Persistence iRules. If using APM <u>prior to version 11.4</u> , enable the built-in _sys_APM_ExchangeSupport_OA_BasicAuth Rule (or if using 11.3.x and NTLM, _sys_APM_ExchangeSupport_OA_NTLMAuth). Important: The Append iRule must be listed first	
	Default Pool	Do <i>not</i> select a default pool for this virtual	
Port 80 (optional, for redirect	Destination Address	IP address for the virtual server (Service Port 80)	
	Profiles	HTTP profile only	
purposes only)	iRule	_sys_https_redirect	

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

This completes the combined virtual server manual configuration table. Continue with <u>Configuration tables for RPC Client Access</u>, <u>POP3</u>, and <u>IMAP4 on page 75</u>.

² Server SSL profile is only necessary if configuring SSL Bridging.

³ If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Auto Map. You must also create the iRule in <u>Oreating an iRule when using a SNAT pool on page 85</u>. See the BIG-IP documentation for creating SNAT Pools. This field is called "Secure Address Translation in version" 11.3 and later.

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

Use this section if you are planning to deploy the BIG-IP system with separate virtual servers for the Exchange CAS services.

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

Outlook Woo App Comit	gai ation table "includes	stile Exchange Conti of Faner (EGF)		
Health Monitors (Main to	ab > Local Traffic > Monitor	rs)		
Follow the monitor guidance for OWA in the table Outlook Web App monitor (includes ECP) on page 69				
Pools (Main tab>Local	Traffic>Pools)			
Health monitor	Add the health monitor you created			
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 (443 if using SSL Bridging) Repeat Address and Port for all members			
iRules (Main tab > Local		9		
iRules	OWA Redirect iRule	Create the Redirect iRule, using the Definition found on page 80		
(Local Traffic>iRules)	Persistence iRule	Create the Persistence iRule, using the Definition found on page 80		
Profiles (Main tab > Loca	al Traffic > Profiles)			
НТТР	Parent Profile	http		
(Profiles>Services)	Redirect Rewrite	All		
HTTP Compression (Profiles>Services)	Content List>Include List	See HTTP Compression Content include list on page 85		
Web Acceleration	Parent Profile	optimized-caching		
(Profiles>Services)	URI List	Add the following to the Exclude list: /owa/ev.owa and uglobal.js		
TCP WAN ³ (Profiles>Protocol)	Parent Profile	tcp-wan-optimized		
TCP LAN³ (Profiles>Protocol)	Parent Profile	tcp-lan-optimized		
Client SSL	Parent Profile	clientssl		
(Profiles>SSL)	Certificate/Key	Select the Certificate and Key you imported		
Server SSL⁴	Parent Profile	serverssl		
(Profiles>SSL)	Options List	If using BIG-IP v11.4.x only, enable No TLSv1.2		
Persistence	Persistence Type	Cookie (Exchange 2010 only)		
OneConnect	Parent Profile	oneconnect		
UneConnect	Source Mask	255.255.255.255		
NTLM	Parent Profile	ntlm		
Virtual Servers (Main tal	b > Local Traffic > Virtual S	ervers)		
	Destination Address	IP address for the virtual server (Service Port 443)		
	Profiles	Add each of the profiles you created above from the appropriate list		
Port 443	SNAT Pool ⁵	Auto Map ⁵		
	iRules	Append, Persistence (the Append iRule must be listed first)		
	Default Pool	Select the pool you created for Outlook Web App above		
Port 80	Destination Address	IP address for the virtual server (Service Port 80)		
(optional, for redirect purposes only)	Profiles	HTTP profile only		
	iRule	_sys_https_redirect		

¹ For Advanced Monitors only. Simple monitors only require the Type, Interval, and Timeout. Replace red text with your FQDN. It must be on a single line with a single \rangle \rangle \rangle \rangle.

² 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.

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

⁴ Server SSL profile is only necessary if configuring SSL Bridging

⁵ If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Auto Map. You must also create the iRule in <u>Creating an iRule when using a SNAT pool on page 85</u>. See the BIG-IP documentation for creating SNAT Pools. This field is "Secure Address Translation in version" 11.3 and later.

Outlook Anywhere configuration table (for separate virtual servers) - includes EWS and OAB

Health Monitors (Main to	ab > Local Traffic > Monit	ors)			
Follow the monitor guidance	Follow the monitor guidance for Outlook Anywhere in the table Outlook Anywhere monitor (includes EWS) on page 69				
Pools (Main tab > Local Traffic > Pools)					
Health monitor	Add the health monitor you created				
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 (443 if using SSL Bridging) Repeat Address and Port for all members				
iRules (Main tab > Local	Traffic > iRules)				
OA Persist	If using Exchange 2010, create the Persistence iRule for Outlook Anywhere, using the Definition found on page 84. You must create this iRule before creating the Persistence profile.				
Profiles (Main tab > Loca	al Traffic > Profiles)				
	Parent Profile	http			
HTTP	Redirect Rewrite	Matching			
TOD WANG	Parent Profile	tcp-wan-optimized			
TCP WAN ²	Nagle's Algorithm	Disabled (clear the Enabled check box)			
TCP LAN ²	Parent Profile	tcp-lan-optimized			
Client SSL	Parent Profile	clientssl			
Ciletti 33L	Certificate/Key	Select the Certificate and Key you imported			
Server SSL ³	Parent Profile	serverssl			
Server 33L	Options List	If using BIG-IP v11.4.x only, enable No TLSv1.2			
OneConnect	Parent Profile	oneconnect			
Oneconnect	Source Mask	255.255.255.255			
NTLM	Parent Profile	ntlm			
	Persistence Type	Universal			
Persistence	iRule	Select the OA Persist iRule you created			
(Exchange 2010 only)		Important: If using 2010, RPC Client Access, and OAB/EWS see <u>Adding the persistence profiles if using</u> Exchange 2010, RPC Client Access, and OAB/EWS on page 78			
Virtual Servers (Main tab	o > Local Traffic > Virtual S	Servers)			
	Destination Address	IP address for the virtual server (Service Port 443)			
	Profiles	Add each of the profiles you created above from the appropriate list			
Port 443	SNAT Pool	Auto Map⁴			
PORT 443	iRules	If using APM <u>prior to version 11.4</u> , enable the built-in <u>sys_APM_ExchangeSupport_OA_BasicAuth</u> rule (or if using 11.3.x and NTLM, <u>sys_APM_ExchangeSupport_OA_NTLMAuth</u>).			
	Default Pool	Select the pool you created for Outlook Anywhere above			
Port 80	Destination Address	IP address for the virtual server (Service Port 80)			
(optional, for redirect purposes only)	Profiles	HTTP profile only			
	iRule	_sys_https_redirect			

For Advanced Monitors only. Simple monitors only require the Type, Interval, and Timeout.
 The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent
 Server SSL profile is only necessary if configuring SSL Bridging.

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

Health Monitors (Main	tab > Local Traffic > Mo	onitors)	
Follow the monitor guidance for Outlook Anywhere in the table ActiveSync monitor on page 69			
Pools (Main tab > Local	Traffic > Pools)		
Health monitor	Add health monitor above	е	
Slow Ramp Time	300		
Load Balancing Method	Least Connections (me	ember) recommended	
Address	IP Address of Client Acce	ess server running ActiveSync	
Service Port	80 (443 if configuring SS	L Bridging) Repeat Address and Port for all members	
Profiles (Main tab > Loc	cal Traffic > Pools)		
HTTP	Parent Profile	http	
TCP WAN ²	Parent Profile	tcp-wan-optimized	
TCP LAN ²	Parent Profile	tcp-lan-optimized	
01:+ 001	Parent Profile	clientssl	
Client SSL	Certificate/Key	Select the Certificate and Key you imported	
Server SSL ³	Parent Profile	serverssl	
(Profiles>SSL)	Options List	If using BIG-IP v11.4.x only, enable No TLSv1.2	
Persistence	Persistence Type	Source Address Affinity (Exchange 2010 only)	
iRules (Main tab > Loca	l Traffic > iRules)		
ActiveSync Persist	If using Exchange 2010,	you can optionally create the Persistence iRule for ActiveSync, using the Definition found on page 80.	
Virtual Servers (Main ta	ab > Local Traffic > Virtu	ual Servers)	
	Destination Address	IP address for the virtual server (Service Port 443)	
	Profiles	Add each of the profiles you created above from the appropriate list	
	SNAT Pool⁴	Auto Map⁴	
Port 443	iRules	Enable the ActiveSync Persist iRule you created. If using APM <u>prior to version 11.4</u> , enable the built-in <u>_sys_APM_ExchangeSupport_OA_BasicAuth</u> rule (or if using 11.3.x and NTLM, <u>_sys_APM_ExchangeSupport_OA_NTLMAuth</u>).	
	Default Pool	Select the pool you created for ActiveSync above	
Port 80	Destination Address	IP address for the virtual server (Service Port 80)	
(optional, for redirect	Profiles	HTTP profile only	
purposes only)	iRule	_sys_https_redirect	

¹ For Advanced Monitors only. Simple monitors only require the Type, Interval, and Timeout. Replace red text with your FQDN. It must be on a single line with a slingle \\n`n.

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 If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Auto Map. You must also create the iRule in <u>Creating an iRule</u> when using a SNAT pool on page 85. See the BIG-IP documentation for creating SNAT Pools. This field is called "Secure Address Translation in version" 11.3 and later.

<u>Autodiscover</u> manual configuration table (for separate virtual server configuration)

Health Monitors (Main tab > Local Traffic > Monitors)		
Follow the monitor guidance for Outlook Anywhere in the table <u>Autodiscover monitor on page 69</u>		
Pools (Main tab > Local	Traffic > Pools)	
Health monitor	Add health monitor above	
Slow Ramp Time	300	
Load Balancing Method	Least Connections (me	mber) recommended
Address	IP Address of Client Acce	ss server running ActiveSync
Service Port	80 (443 if configuring SSI	Bridging) Repeat Address and Port for all members
Profiles (Main tab > Loc	al Traffic > 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
Ollent SSL	Certificate/Key	Select the Certificate and Key you imported
Server SSL ³	Parent Profile	serverssl
(Profiles>SSL)	Options List	If using BIG-IP v11.4.x only, enable No TLSv1.2
Virtual Servers (Main ta	b > Local Traffic > Virtua	al Servers)
	Destination Address	IP address for the virtual server (Service Port 443)
	Profiles	Add each of the profiles you created above from the appropriate list
Port 443	SNAT Pool ⁴	Auto Map⁴
7 071 440	iRules	If using APM <u>prior to version 11.4</u> , enable the built-in _sys_APM_ExchangeSupport_OA_BasicAuth rule (or if using 11.3.x and NTLM, _sys_APM_ExchangeSupport_OA_NTLMAuth)
	Default Pool	Select the pool you created for Autodiscover above
Port 80	Destination Address	IP address for the virtual server (Service Port 80)
(optional, for redirect	Profiles	HTTP profile only
purposes only)	iRule	_sys_https_redirect

¹ For Advanced Monitors only. Simple monitors only require the Type, Interval, and Timeout.

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 If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Auto Map. You must also create the iRule in <u>Creating an iRule</u> when using a SNAT pool on page 85. See the BIG-IP documentation for creating SNAT Pools. This field is called "Secure Address Translation in version" 11.3 and later.

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

(i) Important Exchange Server 2013 Client Access Servers do not offer MAPI as a connection option. If you are deploying Exchange Server 2013, do NOT configure the BIG-IP system for RPC Client Access.

RPC Client Access¹ dynamic port range manual configuration table

Health Monitors (Main ta	b > Local Traffic > Monitors)		
Туре	TCP		
Interval	30 (recommended)		
Timeout	91 (recommended)		
Alias Service Port	135		
Pools (Main tab > Local Ti	raffic > 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 ser	rver running RPC Client Access	
Service Port	* All Services (repeat Address	and Port for all members)	
Profiles (Main tab > Local	l Traffic > Profiles)		
	Parent Profile	Source Address Affinity	
Persistence	Timeout	7200	
(Exchange 2010 only)	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	
	Parent Profile	tcp-wan-optimized	
TCP WAN ³	Idle Timeout	7200	
	Nagle's Algorithm	Disabled (clear the Enabled check box)	
TCP LAN ³	Parent Profile	tcp-lan-optimized	
TCP LAN	Idle Timeout	7200	
Virtual Servers (Main tab	> Local Traffic > Virtual Server	rs)	
	Destination Address	IP address for the virtual server	
	Service Port	135	
Port 135	Profiles	Add each of the profiles you created above from the appropriate list	
	SNAT Pool	Auto Map ⁴	
	Default Pool	Select the pool you created for RPC Client Access above	
	Destination Address	Same IP address used above (make sure you use a unique name)	
All Ports	Service Port	*All Ports	
	Profiles	Add each of the profiles you created above from the appropriate list	
	SNAT Pool	Auto Map ⁴	
	Default Pool	Select the pool you created for RPC Client Access above	
Additional Steps			
After completing this virtual s	erver, you must modify either the S	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.

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.

You must select Advanced from the Configuration list for this option to appear
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.

If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Auto Map. You must also create the iRule in <u>Creating an iRule when using a SNAT pool on page 85</u>. See the BIG-IP documentation for creating SNAT Pools.

RPC Client Access¹ static ports configuration table

Health Monitors (Main ta	ab > Local Traffic > Monitors)		
RPC Monitor			
Туре	TCP		
Interval	30 (recommended)		
Timeout	91 (recommended)		
MAPI Monitor			
Туре	TCP		
Interval	30 (recommended)		
Timeout	91 (recommended)		
Alias Service Port ²	59532 Modify this port to match the RPC Client Access static port for MAPI on your Client Access Servers.		
Address Book Monitor			
Туре	TCP		
Interval	30 (recommended)		
Timeout	91 (recommended)		
Alias Service Port ²	59533 Modify this port to match the RPC Client Access static port for Address Book on your CAS Servers.		
Pools (Main tab > Local 7	Traffic > Pools)		
Health monitor	Add all three health monitors above.		
Availability Requirement	All		
Action on Service Down ²	Reject		
Slow Ramp Time ²	300		
Load Balancing Method	Least Connections (member) recommended		

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 (Main tab > Local Traffic > Profiles)

Address

Service Port

	Parent Profile	Source Address Affinity
Persistence	Timeout	7200
(Exchange 2010 only)	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
TOP WAIN	Idle Timeout	7200
TCP LAN ³	Parent Profile	tcp-lan-optimized
TOP LAIN	Idle Timeout	7200

Virtual Servers (Main tab > Local Traffic > Virtual Servers)

Destination Address IP address for the virtual server

Service Port

Profiles Add each of the profiles you created above from the appropriate list

SNAT Pool Auto Map⁴

Default Pool Select the pool with members using Service Port 135 you created for RPC Client Access above

IP Address of Client Access server running RPC Client Access

135 (repeat Address and Port for all members)

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.

Pools (Main tab > Local Traffic > Pools)

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.

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.

You must select Advanced from the Configuration list for this option to appear

The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent
If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Auto Map. You must also create the iRule in <u>Creating an iRule when using a SNAT pool on page 85</u>. See the BIG-IP documentation for creating SNAT Pools.

POP3 manual configuration table

Health Monitors (Main tab > Local Traffic > Monitors)				
Туре	POP3			
Interval	30 (recommended)	30 (recommended)		
Timeout	91 (recommended)			
User Name	If offloading SSL, typ	e a user name of a POP3 account		
Password	If offloading SSL, typ	e the associated password		
Advanced monitor for PO	P3S (only necessary	if using SSL Bridging)		
Туре	External			
Interval	30 (recommended)			
Timeout	91 (recommended)			
External Program	See Importing the mo	onitor script files on page 79 for the EAV script		
Variables	Name	Value		
	USER PASSWORD DOMAIN EMAIL	The account name associated with a mailbox. The password for the account The Windows domain for the account The email address for the user mailbox (such as j.smith@example.com)		
Pools (Main tab > Local Tr	Pools (Main tab > Local Traffic > Pools)			
Health monitor	Add health monitor a	Add health monitor above		
Slow Ramp Time ¹	300			
Load Balancing Method	Least Connections	Least Connections (member) recommended		
Address	IP Address of Client A	IP Address of Client Access server running POP3		
Service Port	If offloading SSL (PO	If offloading SSL (POP3): 110 If using SSL Bridging (POP3S): 995 (repeat Address and Port for all members)		
Profiles (Main tab > Local	Traffic > Profiles)			
Client SSL	Parent Profile	clientssl		
Olletti GGE	Certificate/Key	Select the Certificate and Key you imported		
Server SSL ³	Parent Profile	serverssl		
(Profiles>SSL)	Options List	If using BIG-IP v11.4.x only, enable No TLSv1.2		
TCP WAN ²	Parent Profile	tcp-wan-optimized		
TCP LAN ²	Parent Profile	tcp-lan-optimized		
Virtual Servers (Main tab > Local Traffic > Profiles)				
Destination Address	IP address for the virtual server			
Service Port	If offloading SSL (POP3): 110 If using SSL Bridging (POP3S): 995			
Profiles	Add each of the profiles you created above from the appropriate list			
SNAT Pool	Auto Map⁴			
Default Pool	Select the pool you o	Select the pool you created for POP3		

You must select Advanced from the Configuration list for this option to appear
 The optimized TCP profiles are optional. If not creating the optimized profiles, create a TCP profile with the base TCP parent

Server SSL profile is only necessary if configuring SSL Bridging.
 If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Auto Map. You must also create the iRule in <u>Creating an iRule when using a SNAT pool on page 85</u>. See the BIG-IP documentation for creating SNAT Pools.

IMAP4 manual configuration table

Health Monitors (Main ta	Health Monitors (Main tab > Local Traffic > Monitors)			
Туре	IMAP4	IMAP4		
Interval	30 (recommended)	30 (recommended)		
Timeout	91 (recommended)			
User Name	If offloading SSL, ty	pe a user name of a IMAP4 account		
Password	If offloading SSL, ty	pe the associated password		
Advanced monitor for IMA	AP4S (only necessar	y if using SSL Bridging)		
Туре	External			
Interval	30 (recommended)			
Timeout	91 (recommended)			
External Program	See Importing the m	onitor script files on page 79 for the EAV script		
Variables	Name	Value		
	USER PASSWORD DOMAIN EMAIL	The account name associated with a mailbox. The password for the account The Windows domain for the account The email address for the user mailbox (such as j.smith@example.com)		
Pools (Main tab > Local Traffic > Pools)				
Health monitor	Add health monitor a	Add health monitor above		
Slow Ramp Time ¹	300			
Load Balancing Method	Least Connections	s (member) recommended		
Address	IP Address of Client	IP Address of Client Access server running IMAP4		
Service Port	If offloading SSL (IM	AP4): 143 If using SSL Bridging (IMAP4S): 993 (repeat Address and Port for all members)		
Profiles (Main tab > Local	l Traffic > Profiles)			
Client SSL	Parent Profile	clientssl		
Cliefft 33L	Certificate/Key	Select the Certificate and Key you imported		
Server SSL ³	Parent Profile	serverssl		
(Profiles>SSL)	Options List	If using BIG-IP v11.4.x only, enable No TLSv1.2		
TCP WAN ²	Parent Profile	tcp-wan-optimized		
TCP LAN ²	Parent Profile	tcp-lan-optimized		
Virtual Servers (Main tab	> Local Traffic > Pro	files)		
Destination Address	IP address for the vi	IP address for the virtual server		
Service Port	If offloading SSL (IM	If offloading SSL (IMAP4): 143 If using SSL Bridging (IMAP4S): 993		
Profiles	Add each of the profiles you created above from the appropriate list			
SNAT Pool	Auto Map ⁴	Auto Map ⁴		
Default Pool	Select the pool you	Select the pool you created for IMAP4		

- 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.

Adding the persistence profiles if using Exchange 2010, RPC Client Access, and OAB/EWS

If you are using Exchange 2010, RPC Client Access, OAB, and EWS, to ensure an even distribution of client connections, you must assign a source address profile as the non-default persistence on the virtual server. This must be performed from the tmsh command line.

Before running the command, you must have first created the virtual server, and the persistence profiles. For the combined virtual server scenario, you create two profiles using default settings: cookie and source address affinity. For the separate Outlook Anywhere virtual, you create the Universal profile described in the table, and a default profile using the source address affinity parent.

Open a command prompt and then type tmsh. From the tmsh prompt, type the following command, replacing <iapp-name> with the name you gave your Exchange iApp:

modify ltm virtual <virtual server name> persist replace-all-with { <cookie persistence or universal
persistence profile name> { default yes } <source IP persistence name> { default no } }

⁴ If you expect more than 6,000 concurrent users per Client Access Server, create a SNAT Pool instead of using Auto Map. You must also create the iRule in <u>Creating an iRule when using a SNAT pool on page 85</u>. See the BIG-IP documentation for creating SNAT Pools.

Monitor script files

This section contains the EAV script and 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.

Importing the monitor script files

Before you can create the advanced monitors for ActiveSync, Autodiscover, POP3S, and/or IMAP4Syou must download and import the applicable monitor files onto the BIG-IP system.

→ Note

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 synchronized.

If you are going to use two instances of the health check to monitor two mail boxes, you must use a unique user name and password for each monitor.

To download and install the script

- Download the appropriate script:
 - Outlook Anywhere (including EWS)
 - » If you configured SSL Offload: http://www.f5.com/pdf/deployment-guides/outlookanywhere-eav-offload.zip
 - » If you configured SSL Bridging http://www.f5.com/pdf/deployment-guides/outlookanywhere-eav-ssl-bridging.zip
 - Autodiscover
 - » If you configured SSL Offload: http://www.f5.com/pdf/deployment-guides/autodiscover-eav-offload.zip
 - » If you configured SSL Bridging http://www.f5.com/pdf/deployment-guides/autodiscover-eav-ssl-bridging.zip
 - POP3S (only necessary if using SSL Bridging and you want to use advanced monitors for POP3S.
 - » http://www.f5.com/pdf/deployment-guides/pop3s-eav.zip
 - IMAP4S (only necessary if using SSL Bridging and you want to use advanced monitors for IMAP4S.
 - » http://www.f5.com/pdf/deployment-guides/imap4s-eav.zip
- Extract the appropriate file(s) to a location accessible by the BIG-IP system.
- 3. From the Main tab of the BIG-IP Configuration utility, expand System, and then click File Management.
- 4. On the Menu bar, click External Monitor Program File List.
- 5. Click the **Import** button.
- In the File Name row, click Browse, and then locate the appropriate file.
- 7. In the **Name** box, type a name for the file related to the script you are using.
- Click the **Import** button.

Now when you create the advanced monitors, you can select the name of the file you imported from the External Program list.

iRules

This section contains the iRules referenced from the manual configuration tables. To create an iRule, from the Main tab, expand **Local Traffic**, and then click **iRules**. Click **Create**, give the iRule a unique name, and then copy and paste the iRule code into the **Definition** section (omitting the line numbers). If specified, you must replace any parts of the code in red text with the names of the appropriate BIG-IP object.

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

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

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

ActiveSync persist iRule

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

If you are using Exchange 2013, do NOT create this iRule.

```
when HTTP_REQUEST {
    if { [HTTP::header exists "Authorization"] } {
        set as_key [sha256 [HTTP::header "Authorization"]]
        persist uie $as_key 7200
} else {
        persist source_addr
}
```

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

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



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

Because of the length of this iRule, you can use the following text file to make the copy paste operation easier: http://www.f5.com/pdf/deployment-guides/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.

If you are using **Exchange 2013**, you must use the iRule in <u>Exchange 2013 only: Persistence iRule if using a single virtual server for all HTTP-based services on page 83</u>.

i Important If you are using Exchange 2010 and RPC Client Access/MAPI, you must uncomment line 69 in the following iRule, as described in line 68.

Exchange 2010 only: Persistence iRule if using a single virtual server for all HTTP-based services

```
## 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
    ## for persistence (if that header is present); otherwise it will
4
5
    ## set persistence according to traditional methods.
6
    ## CHANGE ALL POOL NAMES TO MATCH THOSE IN YOUR ENVIRONMENT.
7
8
    when HTTP_REQUEST {
9
10
11
        ## Offline Address Book and Autodiscover do not require persistence.
12
13
        switch -glob -- [string tolower [HTTP::path]] {
14
15
           "/microsoft-server-activesync*" {
16
              ## ActiveSync.
17
              if { [HTTP::header exists "APM_session"] } {
18
                 persist uie [HTTP::header "APM_session"] 7200
19
              } elseif { [HTTP::header exists "Authorization"] && [string tolower [HTTP::header "Authorization"]] starts_with "basic" } {
20
                  set as_key [sha256 [HTTP::header "Authorization"]]
21
                  persist uie $as_key 7200
22
              } else {
23
                  persist source_addr
24
25
              pool as pool name
26
              COMPRESS::disable
27
              CACHE::disable
28
              return
           }
29
30
           "/owa*" {
31
              ## Outlook Web Access
32
33
                 if { [HTTP::header exists "APM_session"] } {
34
                    persist uie [HTTP::header "APM_session"] 7200
35
                 } else {
36
                    persist cookie insert timeout 0
37
38
                 pool owa_pool_name
39
                 return
40
           }
41
42
           "/ecp*" {
43
              ## Exchange Control Panel.
                 if { [HTTP::header exists "APM_session"] } {
44
45
                    persist uie [HTTP::header "APM_session"] 7200
46
                 } else {
                    persist cookie insert timeout 0
47
48
49
                 pool owa_pool_name
50
                 return
           }
51
52
53
           "/ews*" {
54
              ## Exchange Web Services.
                 if { [HTTP::header exists "APM session"] } {
55
56
                     persist uie [HTTP::header "APM_session"] 7200
57
                 } else {
58
                     persist source_addr
59
60
                 pool oa_pool_name
61
                 COMPRESS::disable
62
                 CACHE::disable
63
                 return
```

81

Critical This iRule continues on the following page.

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

```
65
           "/oab*" {
66
              ## Offline Address Book.
67
                 pool oa_pool_name
68
                 ## uncomment the following line if using RPC Client Access - MAPI
69
                 # persist source_addr
70
                 return
71
           }
72
73
           "/rpc/rpcproxy.dll*" {
74
              ## Outlook Anywhere.
75
              if { [HTTP::header exists "APM_session"] } {
76
                 persist uie [HTTP::header "APM_session"] 7200
77
              } elseif { [HTTP::header exists "Authorization"] && [string tolower [HTTP::header "Authorization"]] starts_with "basic" } {
78
                 set oa_key [sha256 [HTTP::header "Authorization"]]
79
                 persist uie $oa_key 7200
80
              } else {
81
                 persist source_addr
82
83
              pool oa_pool_name
84
              COMPRESS::disable
              CACHE::disable
85
              return
86
87
              }
88
89
           "/autodiscover*" {
90
              ## Autodiscover.
91
              pool ad_pool_name
92
              return
93
           }
94
95
           default {
96
              ## This final section takes all traffic that has not otherwise
97
              ## been accounted for and sends it to the pool for Outlook Web App
98
              if { [HTTP::header exists "APM session"] } {
99
                 persist uie [HTTP::header "APM_session"] 7200
                 } else {
100
                     persist source_addr
101
102
103
                 pool owa_pool_name
104
105
        }
106
    }
107
108
     when HTTP_RESPONSE {
109
        if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate"} {
           ONECONNECT::reuse disable
110
111
           ONECONNECT::detach disable
           ## this command disables NTLM conn pool for connections where OneConnect has been disabled
112
113
           NTLM::disable
114
115
        ## this command rechunks encoded responses
116
        if {[HTTP::header exists "Transfer-Encoding"]} {
117
           HTTP::payload rechunk
118
        }
119 }
```

Exchange 2013 only: Persistence iRule if using a single virtual server for all HTTP-based services

```
## iRule to select pool when all Exchange Client Access HTTP-based services are
    ## accessed through the same BIG-IP virtual server. This iRule is for users who
    ## do not require any persistence.
    ## CHANGE ALL POOL NAMES TO MATCH THOSE IN YOUR ENVIRONMENT.
4
5
6
    when HTTP_REQUEST {
        switch -glob -- [string tolower [HTTP::path]] {
7
8
             "/microsoft-server-activesync*" {
9
                 ## ActiveSync.
10
                 pool as pool name
11
                 COMPRESS::disable
                 CACHE::disable
12
13
                 return
14
             }
15
             "/owa*" {
16
                 ## Outlook Web Access
17
                 pool owa_pool_name
18
                 return
19
20
             "/ecp*" {
21
                 ## Exchange Control Panel.
                 pool owa_pool_name
22
23
24
25
             "/ews*" {
26
                 ## Exchange Web Services.
27
                 pool oa pool name
28
                 COMPRESS::disable
                 CACHE::disable
29
30
                 return
31
             "/oab*" {
32
33
                 ## Offline Address Book.
34
                 pool oa_pool_name
35
                 return
36
37
             "/rpc/rpcproxy.dll*" {
38
                 ## Outlook Anywhere.
39
                 pool oa_pool_name
                 COMPRESS::disable
40
41
                 CACHE::disable
42
                 return
43
             "/autodiscover*" {
44
45
                 ## Autodiscover.
46
                 pool ad_pool_name
47
                 return
48
49
             default {
50
                ## This final section takes all traffic that has not otherwise
51
                ## been accounted for and sends it to the pool for Outlook Web App
52
               pool owa pool name
53
             }
54
        }
55
    when HTTP_RESPONSE {
56
57
        if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate"} {
58
             ONECONNECT::reuse disable
59
             ONECONNECT::detach disable
60
             ## this command disables NTLM conn pool for connections where OneConnect
61
             ## has been disabled
62
             NTLM::disable
63
64
        ## this command rechunks encoded responses
65
           {[HTTP::header exists "Transfer-Encoding"]} {
66
             HTTP::payload rechunk
67
68
    }
```

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. Use the appropriate iRule depending on your version of Exchange.

Outlook Anywhere persistence iRule for Exchange 2010 only

```
when HTTP REQUEST {
       switch -glob -- [string tolower [HTTP::path]] {
2
3
           "/ews*" {
4
          ## Exchange Web Services.
5
              if { [HTTP::header exists "APM_session"] } {
                 persist uie [HTTP::header "APM session"] 7200
6
7
                 } else {
8
                 persist source_addr
9
10
           "/rpc/rpcproxy.dll*" {
11
12
           ## Outlook Anywhere.
              if { [HTTP::header exists "APM_session"] } {
13
14
                 persist uie [HTTP::header "APM_session"] 7200
                 } elseif { [string tolower [HTTP::header "Authorization"]] starts_with "basic" } {
15
16
                 set oa_key [sha256 [HTTP::header "Authorization"]]
17
                 persist uie $oa_key 7200
18
                 } else {
19
                 persist source addr
20
21
          }
22
23
    }
    when HTTP_RESPONSE {
24
        if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate"} {
25
             ONECONNECT::reuse disable
26
27
             ONECONNECT::detach disable
             ## disables NTLM conn pool for connections where OneConnect has been disabled
28
29
             NTLM::disable
30
31
        ## this command rechunks encoded responses
32
        if {[HTTP::header exists "Transfer-Encoding"]} {
33
            HTTP::payload rechunk
34
        }
35
    }
```

Outlook Anywhere persistence iRule for Exchange 2013 or deployments not using persistence only

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

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 (omitting the line numbers) in the **Definition** section. You need one IP address for each 6,000 concurrent users you expect to each Client Access Server. Modify the IP addresses in the following example to your SNAT Pool IP addresses, adding or removing lines as necessary.

Make sure to attach the iRule to the virtual servers where you are using a SNAT pool.

```
when RULE_INIT {
2
       # Use a local array to configure SNAT addresses.
3
        # These addresses must be defined in a SNAT pool.
4
        # In this example, we use three addresses. Replace
5
        # these with the IP addresses used in your SNAT Pool.
        # Follow the pattern of the existing addresses to add more than three.
6
8
        set static::snat_ips(0) 10.0.0.1
9
        set static::snat_ips(1) 10.0.0.2
10
       set static::snat_ips(2) 10.0.0.3
11
12
     when CLIENT_ACCEPTED {
13
14
       # Calculate the crc32 checksum of the client IP.
15
       # Use the modulo of the checksum and the number of SNAT IPs in the array
       # to select a SNAT IP address.
16
17
18
        snat $static::snat_ips([expr {[crc32 [IP::client_addr]] % [array size static::snat_ips]}])
19
     }
20
```

Note: If you are configuring multiple Exchange deployments on the same BIG-IP device and are using SNAT pools, you must change the variable names (snat_ips) in the iRule for each separate deployment.

HTTP Compression Content include list

Use the following list for the Content list in the HTTP Compression profiles

- text/(css | html | javascript | json | plain | postscript | richtext | rtf | vnd.wap.wml | vnd.wap.wmlscript | wap | wml | x-component | x-vcalendar | x-vcard | xml)
- application/(css | css-stylesheet | doc | excel | javascript | json | lotus123 | mdb | mpp | ms-excel | ms-powerpoint | ms-word | msaccess | msexcel | mspowerpoint | msproject | msword | photoshop | postscript | powerpoint | ps | psd | quarkexpress | rtf | txt | visio | vnd.excel | vnd.ms-access | vnd.ms-excel | vnd.ms-powerpoint | vnd.ms-pps | vnd.ms-project | vnd.ms-word | vnd.ms-works | vnd.ms-works-db | vnd.msaccess | vnd.msexcel | vnd.mspowerpoint | vnd.msword | vnd.msword | vnd.powerpoint | vnd.visio | vnd.wap.cmlscriptc | vnd.wap.wmlc | vnd.wap.xhtml+xml | vnd.word | vsd | winword | wks | word | x-excel | x-java-jnlp-file | x-javascript | x-json | x-lotus123 | x-mdb | x-ms-excel | x-ms-project | x-mscardfile | x-msclip | x-msexcel | x-mspowerpoint | x-msproject | x-msword | x-msworks-db | x-msworks-wps | x-photoshop | x-postscript | x-powerpoint | x-ps | x-quark-express | x-rtf | x-vermeer-rpc | x-visio | x-vsd | x-wks | x-word | x-xls | x-xml | xhtml+xml | xls | xml)
- image/(photoshop | psd | x-photoshop | x-vsd)

BIG-IP APM manual configuration

This section covers the following scenarios for BIG-IP APM:

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



Figure 4: BIG-IP APM with SSL Offload configuration example

b. Both the BIG-IP APM 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.



Figure 5: BIG-IP APM with SSL Offload configuration example

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

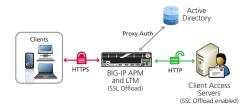


Figure 6: BIG-IP APM with SSL Bridging configuration example

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

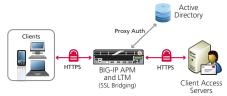


Figure 7: BIG-IP APM with SSL Bridging configuration example

BIG-IP APM Configuration

No matter which of the scenarios you are deploying, use the following table to create the BIG-IP APM configuration (scenario-specific configuration begins after this section). 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.

Powershell command for enabling the OWA logon options

If you want to display the computer type (public/shared vs private) and light version ("Use the light version of Outlook Web App") options for Outlook Web App on the APM logon page via the BIG-IP APM, you must run the following PowerShell command on one of your Client Access Servers (only one):

Get-OwaVirtualDirectory | Set-OwaVirtualDirectory -LogonPageLightSelectionEnabled \$true -LogonPagePublicPrivateSelectionEnabled \$true

Give each BIG-IP object a unique name in the Name field.

DNS and NTP

See Configuring DNS and NTP settings on page 65 for instructions.

Health Monitors (Main tab > Local Traffic > Monitors)

Configuration Select Advanced from the Configuration list (if necessary).

Type LDAP

Interval10 (recommended)Timeout31 (recommended)

User Name Type a user name with administrative permissions

Password Type the associated password

Base Specify your LDAP base tree. For example, CN=Exchange Users, DC=example, DC=com

Filter Specify the filter. We type cn=user1, using the example above: user1 in OU group "Exchange Users" and domain "example.com"

Security Select a Security option (either None, SSL, or TLS)

Chase Referrals Yes

Alias Address *All Addresses

Alias Address Port 389 (for None or TLS) or 636 (for SSL)

AAA Server (Main tab-->Access Policy-->AAA Servers)

Type Active Directory

Domain Name Type the FQDN of the Windows Domain name

Server Connection Click Use Pool if necessary.

Domain Controller Pool Name Type a unique name

Domain Controllers IP Address: Type the IP address of a domain controller

Hostname: Type the FQDN of the domain controller

 $\label{eq:click} \textbf{Add}. \ \ \text{Repeat for each domain controller in this configuration}.$

Server Pool Monitor
Select the monitor you created above.

Admin Name!
Type the Administrator name

Admin Name¹Type the Administrator nameAdmin Password¹Type the associated password

SSO Configuration² (Main tab-->Access Policy-->SSO Configurations)

Forms based SSO Configuration

SSO Configurations By Type Forms-Client Initiated

SSO Configuration Name Type a unique name. We use Exchange-SSOv2

In the left pane of the box, click Form Settings, and then click Create.

Form Name Type a unique name. We use Exchange-Form

¹ Optional; Admin Name and Password are only required if anonymous binding to Active Directory is not allowed in your environment

² If you are using BIG-IP version 11.3, you can optionally create a Kerberos SSO configuration for Outlook Anywhere. See <u>Optional: Configuring the APM for Outlook Anywhere with NTLM Authentication - BIG-IP v11.3 or later only on page 102</u>

Forms based SSO Configuration continued

In the left pane of the box, click Form Parameters, and then click Create

Form Parameter Name Select Username from the list.

Username Parameter Value %{session.sso.token.last.username}

Click Ok, and then click Create again in the Forms Parameters box.

Form Parameter Name Select password

Form Parameter Value Select %{session.sso.token.last.password}

Secure Yes

Form Parameters Click Ok. If you are not using the OWA logon options, continue with Form Detection. If

you are using OWA logon options, click **Create** again in the Forms Parameters box.

Form Parameter Name Type flags

Form Parameter Value Type %{session.custom.owa.flags}
Click Ok, and then click Create again in the Forms Parameters box.

Form Parameter Name Type trusted

Form Parameter Value Type %{session.custom.owa.trusted}

Form Detection In the left page of the Create New Form Definition box, click Form Detection.

Detect Form by Select URI

Request URI Type /owa/auth/logon.aspx.

Logon Detection In the left page of the Create New Form Definition box, click Logon Detection.

Detect Logon by Select Presence of Cookie

Cookie Name Type cadata

JavaScript Injection In the left page of the Create New Form Definition box, click JavaScript Detection.

Injection Method Select extra

Extra JavaScript Type clkLgn() Click Ok twice to complete the SSO Configuration.

NTLM SSO Configuration (create only one)

NTLMv1 (create this object if you are using NTLMv1)

SSO Method NTLMv1 (if you are using NTLMv2 only, select NTLMv2)

Username Conversion Enable

NTLM Domain The NTLM domain name where the user accounts are located

NTLMv2 (create this object if you are using only NTLMv2)

SSO Method NTLMv2

NTLM Domain Enter the fully-qualified name of the domain where users will authenticate

11.4 and later only: Exchange Profile3 (Main tab > Access Policy > Application Access > Microsoft Exchange)

Parent Profile /Common/exchange

In the left pane of the box, under Service Settings, click Autodiscover

SSO Configuration From the Autodiscover SSO Configuration list, select the NTLM SSO Configuration you created.

In the left pane of the box, under Service Settings, click Exchange Web Service

SSO Configuration From the EWS SSO Configuration list, select the NTLM SSO Configuration you created.

In the left pane of the box, under Service Settings, click Offline Address Book

SSO Configuration From the OAB SSO Configuration list, select the NTLM SSO Configuration you created.

Access Profile (Main tab > Access Policy > Access Profiles)

Microsoft Exchange⁴ If you created the Exchange profile, select the profile you created from the list.

SSO Configuration If using BIG-IP v11.3 or earlier, select the name of NTLM SSO configuration you created

Edit the Access Policy

Edit the Access Profile you just created using the Visual Policy Editor. Continue now with Editing the Access Policy on page 89

³ If using the Exchange profile in 11.4 and later, you must remove any _sys_APM irules from the virtual server

⁴ Optional, only available in 11.4 and later, and only applicable if you created the Exchange profile.

Editing the Access Policy

After creating the objects in the table above, use the following procedure to edit the Access Policy on the BIG-IP APM using the Visual Policy Editor (VPE). The Policy shown 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, and then, in the Access Policy column, click Edit.
- 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. From the Split domain from full Username list, select Yes.
 - c. Configure the rest of the Logon Page properties as applicable, and then click Save.
- 4. Click the + symbol between Logon Page and Deny.
 - a. In the Authentication section, click the AD Auth option button, and click Add Item.
 - b. In the **Active Directory** properties box, from the **Server** list, select the AAA server you created using the table above. The rest of the settings are optional. Click **Save**.
- 5. On the Successful path between **AD Auth** and **Deny**, click the **+** symbol.
 - a. Click the SSO Credential Mapping option button, and then click Add Item.
 - b. Configure the Properties as applicable for your configuration; we leave the settings at the defaults. Click Save.
- 6. On the fallback path between **SSO Credential Mapping** and **Deny**, click the **Deny** box. Click the **Allow** option button, and then click **Save**. See Figure 8.
- 7. Click the Apply Access Policy link in the upper left part of the window. You must apply an access policy before it takes effect.



Figure 8: Example of the Access Policy in the VPE

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 4. If you are in a partition other than /Common, replace /Common with the name of your partition.

```
when ACCESS_ACL_ALLOWED {
    set req_uri [HTTP::uri]
    if { $req_uri contains "/owa/auth" } {
        WEBSSO::select [set foo /Common/Exchange-SSOv2]
    }
    unset req_uri
}
```

Exchange 2013 only If using Exchange 2013, line 3 is: if { \$req_uri contains "/owa" } {

Click the Finished button.

Configuration table for scenario 1: BIG-IP APM sending traffic to a remote BIG-IP LTM

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

There are additional procedu	ures immediately following	g this table.	
Health Monitors (Main tab >	· Local Traffic > Monitors)		
Туре	TCP		
Interval	30 (recommended)		
Timeout	91 (recommended)		
Pools (Main tab > Local Traffi	,		
Health Monitor	Select the monitor you create	d above	
Load Balancing Method	Round Robin	4 450.0	
Address		e BIG-IP LTM virtual server to which this BIG-IP APM will forward traffic	
Service Port	•	e-encrypting for SSL Bridging	
iRules (Main tab > Local Trafi		Pend ypting for OSE Bridging	
See <u>Creating the persist iRule on OWA (default) on page 94</u> or <u>Option</u>	the BIG-IP APM on page 91 and onal: Creating the iRule to termin	d Creating the iRule to terminate inactive APM sessions if using Forms-based authentication for inate inactive APM sessions if using Windows based authentication for OWA on page 95. You do not be a stopped iRule on page 80.	
Profiles (Main tab > Local Tra	affic > Profiles)		
HTTP (Profiles>Services)	Parent Profile	http	
HTTP Compression	Parent Profile	wan-optimized-compression	
(Profiles>Services)	Content List>Include List	See HTTP Compression Content include list on page 85	
Web Acceleration	Parent Profile	optimized-caching	
(Profiles>Services)	URI List	Add the following to the Exclude list: /owa/ev.owa and uglobal.js	
TCP WAN (Profiles>Protocol)	Parent Profile	tcp-wan-optimized	
TCP LAN (Profiles>Protocol)	Parent Profile	tcp-lan-optimized	
OneConnect (Profiles>Other)	Parent Profile	oneconnect	
NTLM (Profiles>Other)	Parent Profile	ntlm	
Client SSL (Profiles>SSL)	Parent Profile	clientssl	
Cheff GGE (Fromes >GGE)	Certificate and Key	Select your Certificate and key	
Server SSL (for SSL Bridging only)	Parent Profile	If the remote BIG-IP LTM receiving this traffic is using a self-signed or default certificate for decryption, select serverssl-insecure-compatible If it's using a certificate signed by a Certificate Authority, select serverssl	
(Profiles>SSL)	Certificate and Key	Select your Certificate and key	
	Options List	If using BIG-IP v11.4.x only, enable No TLSv1.2	
Virtual Servers (Main tab > L	Local Traffic > Virtual Servers	s)	
Destination Address	The IP address clients use to	access Exchange. Your Exchange FQDN resolves to this IP address.	
Service Port	443		
OneConnect profile	Select the OneConnect profile	e you created	
HTTP Profile	Select the HTTP profile you o	created above	
HTTP Compression Profile	Select the HTTP Compression	n profile you created	
Web Acceleration Profile	Select the Web Acceleration	profile you created	
SSL Profile (Client)	Select the Client SSL profile you created		
SSL Profile (Server)	Select the Server SSL profile	you created (only for Scenario 2, SSL Bridging).	
Access Profile	Select the Access Profile you	created	
iRules¹	Enable the Append iRule you created on page 80. Enable the iRule you created to terminate inactive sessions Enable either the built-in _sys_APM_ExchangeSupport_OA_BasicAuth or sys_APM_ExchangeSupport_OA_NTLMAuth Rule as depending on your auth method. This rule is necessary whether deploying Outlook Anywhere or not.¹ Enable the iRule that chooses the SSO configuration you created (select_SSO_irule in our example) Enable the APM session ID irule you created (apm-irule in our example)		
Default Deal	Colort the Declarate greated		

¹ Do not attach the _sys_APM_ExchangeSupport_OA_BasicAuth iRule if you are using BIG-IP v11.4 <u>and</u> the Exchange profile.

Select the Pool you created

Default Pool

Creating the persist iRule on the BIG-IP APM

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

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

- 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-session-id-irule.
- 3. In the **Definition** section, copy and paste the following iRule, omitting the line numbers.

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

4. Click the **Finished** button.

BIG-IP LTM iRule if all traffic goes through the BIG-IP APM

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

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

To create the persistence iRule if all traffic goes through the BIG-IP APM to the LTM

- 1. On the Main tab, expand Local Traffic, click iRules, and then click the Create button to create a new iRule.
- 2. In the **Name** box, type a unique name. In our example, we type **apm-persist**.
- 3. In the **Definition** section, copy and paste the appropriate iRule (omitting the line numbers), depending on your version of Exchange.

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

Click the Finished button.

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

Modifying the virtual server to use the new persistence iRule

If you just created the new persistence iRule on the BIG-IP LTM (<u>BIG-IP LTM iRule if all traffic goes through the BIG-IP APM on page 91</u>), and have an existing BIG-IP LTM configuration, you must modify the BIG-IP LTM virtual server to use the new persistence iRule and remove any existing persistence iRules.

This completes the BIG-IP APM configuration for scenario 1.

Creating the iRule to terminate inactive APM sessions if using Forms-based authentication for OWA (default)

When using APM to secure OWA, APM sessions can remain active after users have 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. **Note:** This iRule is only effective if you are using Forms-based authentication for OWA, if using Windows Authentication, see Optional: Creating the iRule to terminate inactive APM sessions if using Windows based authentication for OWA on page 95.

To add the APM 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. Note that line 22 is a single line.

```
1
    when RULE INIT {
2
        set static::cookie_sessionid [format "sessionid=null; path=/; Expires=Thurs, 01-Jan-1970 00:00:00 GMT;"]
        set static::cookie_cadata [format "cadata=null; path=/; Expires=Thurs, 01-Jan-1970 00:00:00 GMT;"]
3
        set static::cookie_usercontext [format "UserContext=null; path=/; Expires=Thurs, 01-Jan-1970 00:00:00 GMT;"]
4
5
    }
6
    when ACCESS_SESSION_STARTED {
        if { [string tolower [HTTP::uri]] contains "ua=0" } {
7
8
           ACCESS::session remove
9
10
   }
11
    when ACCESS_ACL_ALLOWED {
12
        set apm mrhsession [HTTP::cookie value "MRHSession"]
        if { [table lookup $apm_mrhsession] == "EXCHANGE_LOGOUT" } {
13
14
           ACCESS::session remove
15
           table delete $apm_mrhsession
16
        }
17
    }
    when HTTP_REQUEST {
18
        set isset 0
19
20
        if {[string tolower [HTTP::uri]] starts_with "/owa" } {
21
           if {[string tolower [HTTP::uri]] contains "logoff" } {
22
               ACCESS::session remove
    HTTP::respond 302 Location "https://[HTTP::host]/vdesk/hangup.php3" "Set-Cookie" $static::cookie_sessionid "Set-Cookie" $static::cookie_sessionid "Set-Cookie" $static::cookie_sessionid "Set-Cookie" $static::cookie_usercontext
23
24
              if { [string tolower [HTTP::uri]] contains "ua=0" } {
25
26
                  set mrhsession [HTTP::cookie value "MRHSession"]
27
                  set isset 1
              }
28
29
           }
30
        }
31
    }
    when HTTP_RESPONSE {
32
33
        if { $isset == 1 } {
           if { $mrhsession != "" && [HTTP::status] == 440 } {
34
35
              table set $apm_mrhsession "EXCHANGE_LOGOUT"
36
               return
37
        }
38
   }
39
```

- Click Finished.
- 6. On the Main tab, click Virtual Servers.
- 7. From the **Virtual Server** list, click the name of the appropriate virtual server (either the BIG-IP APM virtual server, the combined virtual server, or the separate OWA virtual server, depending on how you configured the BIG-IP system.
- 8. On the Menu bar, click Resources.
- 9. From the iRules section, click Manage.
- 10. From the **Available** list, select the iRule you just created and then click Add (<<).
- 11. If deploying for BIG-IP APM, click the **Up** button to move the this iRule just below the **<iapp-name>_sys_APM_ ExchangeSupport_OA_BasicAuth** (or **<iapp-name>_sys_APM_ExchangeSupport_OA_NtImAuth** if using NTLM for OA) rule. If you are using BIG-IP version 11.4 and deploying with the BIG-IP APM Exchange profile, this step is not necessary.
- 12. Click Finished.

Optional: Creating the iRule to terminate inactive APM sessions if using Windows based authentication for OWA

When using APM to secure OWA, APM sessions can remain active after users have closed the OWA window. This iRule checks the OWA session status and terminates the associated APM session after a configurable amount time.

Note: This iRule is only effective if you are using Windows based authentication for OWA.

To modify the APM session check iRule if you used the iApp template

- 1. If you have not already disabled Strict Updates, see Step 2. Disable the Strict Updates feature: on page 55.
- 2. On the Main tab, click iRules.
- 3. From the iRules list, click <iapp-name>_login_timeout.
- 4. Follow the instructions in Step 3 of the next procedure to copy and paste the iRule in the Definition section.
- Click Update.

To add the APM session check iRule if you are configuring the system manually

- 1. On the Main tab, click Local Traffic > iRules > Create.
- 2. In the Name box, type a unique name such as apm-owa-session-irule.
- 3. In the **Definition** section, copy and paste the following iRule. Note that line 15 is a single line. You can also modify the timeout values by changing the values in red in lines 10, 13, and 19. For example, if you wanted to change the timeout value to 15 minutes, change **1800** to **2700**.

```
1
    when RULE INIT {
       set static::cookie_sessionid [format "sessionid=null; path=/; Expires=Thurs, 01-Jan-1970 00:00:00 GMT;"]
2
       set static::cookie_outlooksession [format "OutlookSession=null; path=/; Expires=Thurs, 01-Jan-1970 00:00:00 GMT;"]
3
       set static::cookie_cadata [format "cadata=null; path=/; Expires=Thurs, 01-Jan-1970 00:00:00 GMT;"]
5
       set static::cookie_usercontext [format "UserContext=null; path=/; Expires=Thurs, 01-Jan-1970 00:00:00 GMT;"]
6
    }
7
    when HTTP_REQUEST {
8
       if {[string tolower [HTTP::uri]] starts_with "/owa" } {
9
          set owa session [HTTP::cookie value "OutlookSession"]
10
          table add $owa_session 0 3600
11
          if {[string tolower [HTTP::uri]] contains "ua=0" } {
12
             table incr -notouch $owa_session
13
             if {[table lookup -notouch $owa_session] != 0 && [table timeout -remaining $owa_session] < 1800 } {
14
                log local0. "Session timed out"
15
                HTTP::respond 440 Set-Cookie $static::cookie_sessionid Set-Cookie $static::cookie_outlooksession Set-Cookie
    $static::cookie_cadata Set-Cookie $static::cookie_usercontext
16
                ACCESS::session remove
             }
17
18
          } else {
19
             table replace $owa_session 0 3600
20
21
       }
   }
22
```

- 4. Click Finished.
- 5. The next step depends on whether you used the iApp template or are configuring the BIG-IP system manually
- 6. On the Main tab, click Virtual Servers.
- 7. From the **Virtual Server** list, click the name of the appropriate virtual server (either the BIG-IP APM virtual server, the combined virtual server, or the separate OWA virtual server, depending on how you configured the BIG-IP system.
- 8. On the Menu bar, click **Resources**, and then from the iRules section, click **Manage**.
- 9. From the **Available** list, select the iRule you just created and then click Add (<<).
- 10. If deploying for BIG-IP APM, click the **Up** button to move the this iRule just below the **<iapp-name>_sys_APM_ ExchangeSupport_OA_BasicAuth** (or **<iapp-name>_sys_APM_ExchangeSupport_OA_NtlmAuth** if using NTLM for OA) rule. If you are using BIG-IP version 11.4 and deploying with the BIG-IP APM Exchange profile, this step is not necessary.
- 11. Click Finished.

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

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

Creating the persistence iRule when using BIG-IP APM

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

To create the iRule

- 1. On the Main tab, expand Local Traffic, click iRules, and then click the Create button.
- In the Name box, give the iRule a unique name. We use apm-persistence-irule.
- 3. In the **Definition** section, copy and paste the appropriate iRule, depending on your version of Exchange.

BIG-IP APM iRule for Exchange 2010 only

```
## CHANGE ALL POOL NAMES TO MATCH THOSE IN YOUR ENVIRONMENT.
    ## OAB and Autodiscover do not require persistence.
2
3
    when ACCESS_ACL_ALLOWED {
4
          set sessionid [ACCESS::session data get "session.user.sessionid"]
          switch -glob -- [string tolower [HTTP::path]] {
5
               "/microsoft-server-activesync*" {
6
7
                    pool my_Exchange_2010__single_as_pool
8
                    COMPRESS::disable
9
                    CACHE::disable
10
                    persist uie $sessionid 7200
11
                    return
12
               "/ews*" {
13
14
                    pool my Exchange 2010 single oa pool
                    COMPRESS::disable
15
16
                    CACHE::disable
17
                    persist uie $sessionid 7200
18
                    return
19
               "/ecp*" {
20
                    pool my_Exchange_2010__single_owa_pool
21
22
                    persist uie $sessionid 7200
23
                    return
              }
"/oab*" {
24
25
26
                    pool my_Exchange_2010__single_oa_pool
27
                    return
28
               "/rpc/rpcproxy.dll*" {
29
                    pool my_Exchange_2010__single_oa_pool
30
31
                    COMPRESS::disable
                    CACHE::disable
32
33
                    persist uie $sessionid 7200
34
                    return
35
               "/autodiscover*" {
36
37
                    pool my_Exchange_2010__single_ad_pool
                    return
38
39
               default {
40
41
               ## This final section takes all traffic that has not otherwise
42
               ## been accounted for and sends it to the pool for Outlook Web App
43
                    pool my_Exchange_2010__single_owa_pool
44
                    persist uie $sessionid 7200
45
               }
46
          }
47
    when HTTP RESPONSE {
48
49
         if { [string tolower [HTTP::header values "WWW-Authenticate"]] contains "negotiate"} {
50
             ONECONNECT::reuse disable
51
             ONECONNECT::detach disable
52
             ## disables NTLM conn pool for connections where OneConnect has been disabled
53
             NTLM::disable
54
55
         ## this command rechunks encoded responses
            {[HTTP::header exists "Transfer-Encoding"]} {
56
             HTTP::payload rechunk
57
58
59
    }
```

4. Click **Finished**.

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

Modifying the virtual server to use the iRules and Access Profile

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

If you created separate virtual servers, you must add the persistence iRule and Access Profile to all BIG-IP LTM virtual server for the HTTP-based Client Access Services (Outlook Web App, Outlook Anywhere, ActiveSync, and Autodiscover). The terminate inactive sessions iRule only needs to be assigned to the OWA virtual server.

Optional: Securing Access to the Exchange 2013 Administration Center with BIG-IP APM

In Microsoft Exchange Server 2013, Exchange administration is now performed via a web-based console, the Exchange Administration Center (EAC). You can use F5's APM module to query Active Directory group membership for the user making the request to EAC. If the user is not a member of the Organization Management group, the APM policy denies access.

Creating the Access profile

This configuration requires creating a new APM Access Profile object. If you have previously deployed Exchange 2010 CAS servers with APM using the iApp template, the simplest way is to create the profile is to copy the existing policy created by the template.

Copying the Access Policy created by the iApp template

To copy the Access Policy created by iApp, use the following procedure.

To copy the Access Policy created by the iApp template

- 1. On the Main tab, expand Access Policy, and click Access Profiles.
- 2. In the Access Policy list, find the row for the Access Policy created by the Exchange iApp template. This policy starts with the name you gave the iApp, followed by **_apm_access**.
- 3. Click the Copy link that corresponds to the Access Policy.
- 4. In the **Copied Profile Name** box, type a new name for this profile.
- 5. Click the Copy button.
- 6. Continue with Editing the APM Access Policy if you copied the existing Access Policy on page 100.

Creating a new Access Policy

To create a new Access Policy, use the following table for guidance. For specific instructions, see the online help or product manuals.

BIG-IP APM Object	Non-default settings/Notes	
Access Profile	Name	Type a unique name.
(Main tab>Access Policy>Access Profiles)	SSO Configuration	Use the NTLMv1 SSO object created by the iApp template

Editing the APM Access Policy if you created a new Access Policy

Use this section to edit the Access Profile if you created a new Access Policy.

To edit 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. From the **Split domain from full Username** list, select **Yes**.
 - c. Configure the rest of the Logon Page properties as applicable, and then click Save.
- 4. Click the + symbol between Logon Page and Deny.
 - a. In the Authentication section, click the AD Auth option button, and click Add Item.
 - b. In the **Active Directory** properties box, from the **Server** list, select the AAA server created by the iApp.
 - c. 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 **Empty** option button, and then click **Add Item**.
 - b. In the Name box, type EAC URI Check.

- c. Click the Branch Rules tab.
- d. Click Add Branch Rule.
- e. In the Name box, type is EAC.
- f. In the Expression row, click the **change** link.
- g. Click Add Expression.
- h. From the Agent Sel list, select Landing URI.
- i. In the Landing URI is box, type /ecp/default.aspx.
- j. Click Add Expression.
- k. Click the Finished button.
- I. Click the **Save** button.
- 6. On the **is EAC** path (if you did not modify the name, this is Branch Rule 1) between **EAC URI Check** and **Deny** click the **+** symbol.
 - a. Click AD Query, and then click Add Item.
 - b. In the Name box, type EAC AD Query.
 - c. From the **Server** list, select the AAA server created by the iApp.
 - d. In the Search Filter box, type sAMAccountName=%{session.logon.last.username}.
 - e. Click the Branch Rules tab.
 - f. In the Name box, delete any existing text, and then type Organization Management.
 - g. In the Expression row, click the change link.
 - h. Click the Delete (x) button to the right of the User's Primary Group ID is box.
 - i. Click Add Expression.
 - j. From the Agent Sel list, select AD Query.
 - k. From the Condition list, select User is a Member of.
 - In the User is a member of box, type CN=Organization Management,OU=Microsoft Exchange Security Groups,DC=example,DC=com.
 - m. Click Add Expression.
 - n. Click the Finished button.
 - o. Click the Save button.
- 7. On the fallback path between **EAC URI Check** and **Deny**, click the + symbol.
 - a. Click SSO Credential Mapping, and then click Add Item.
 - b. Configure the settings as applicable. We leave the settings at the defaults.
 - c. Click Save.
 - d. On the fallback path between SSO Credential Mapping and Deny, click the Deny box/link.
 - e. Click the Allow option button, and then click Save.
- 8. On the Organization Management path, between EAC AD Query and Deny click +.
 - a. Click SSO Credential Mapping, and then click Add Item.
 - b. Configure the settings as applicable. We leave the settings at the defaults.
 - c. Click Save.
 - d. On the fallback path between SSO Credential Mapping and Deny, click the Deny box/link.
 - e. Click the **Allow** option button, and then click **Save**.

- 9. 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.
- 10. Continue with Modifying the virtual server to use the new Access Policy on page 101.

When you are finished, your VPE should look like the following:



Figure 9: Example of the Access Policy in the VPE

Editing the APM Access Policy if you copied the existing Access Policy

Use this section to edit the Access Profile if you made a copy of the Access Policy created by the iApp template.

To edit the access policy

- 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. On the Successful path between AD Auth and SSO Credential Mapping, click the + symbol.
 - a. Click the **Empty** option button, and then click **Add Item**.
 - b. In the Name box, type EAC URI Check.
 - c. Click the Branch Rules tab.
 - d. Click Add Branch Rule.
 - e. In the Name box, type is EAC.
 - f. In the Expression row, click the **change** link.
 - g. Click Add Expression.
 - h. From the Agent Sel list, select Landing URI.
 - i. In the Landing URI is box, type /ecp/default.aspx.
 - j. Click Add Expression.
 - k. Click the Finished button.
 - I. Click the **Save** button.
- 4. On the **is EAC** path (if you did not modify the name, this is Branch Rule 1) between **EAC URI Check** and **Deny** click the + symbol.
 - a. Click AD Query, and then click Add Item.
 - b. In the Name box, type EAC AD Query.
 - c. From the **Server** list, select the AAA Server created by the iApp.
 - d. In the Search Filter box, type sAMAccountName=%{session.logon.last.username}.
 - e. Click the Branch Rules tab.
 - f. In the Name box, delete any existing text, and then type Organization Management.

- g. In the Expression row, click the change link.
- h. Click the Delete (x) button to the right of the User's Primary Group ID is box.
- i. Click Add Expression.
- j. From the Agent Sel list, select AD Query.
- k. From the Condition list, select User is a Member of.
- In the User is a member of box, type CN=Organization Management,OU=Microsoft Exchange Security Groups,DC=example,DC=com.
- m. Click Add Expression.
- n. Click the Finished button.
- o. Click the Save button.
- 5. On the Organization Management path, between EAC AD Query and Deny click the + symbol.
 - a. Click SSO Credential Mapping, and then click Add Item.
 - b. Configure the settings as applicable. We leave the settings at the defaults.
 - c. Click Save.
 - d. On the fallback path between SSO Credential Mapping and Deny, click the Deny box/link.
 - e. Click the Allow option button, and then click Save.
- 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.
- 7. Use the following procedure to add the Access Policy to the virtual server.

Modifying the virtual server to use the new Access Policy

The final task is to add the new Access Policy to the virtual server.

To modify the virtual server to use the Access Policy

- 1. On the Main tab, expand Local Traffic and then click Virtual Servers.
- 2. Click the name of the appropriate virtual server. This is either the single virtual server for all HTTP-based CAS services or the separate virtual server for OWA.
- 3. In the Access Policy section, from the Access Profile list, select the Access profile you just modified.
- 4. Click Update.

This completes the EAC configuration.

Optional: Configuring the APM for Outlook Anywhere with NTLM Authentication - BIG-IP v11.3 or later only

F5's Access Policy Manager (APM) module supports NTLM authentication for Outlook clients using the RPC-over-HTTP protocol (Outlook Anywhere) in version 11.3 and later. Use the following table for guidance on configuring the BIG-IP APM. Give each BIG-IP object a unique name in the Name field.

Before configuring BIG-IP system, you must perform prerequisite configuration steps on the Exchange Server(s) and Active Directory servers. See Appendix E: Active Directory and Exchange Server configuration for NTLM on page 115.

AAA Server (Access Policy:	>AAA Servers)	
Туре	Active Directory	
Domain Controller	Type the IP address or FQDN name of an Active Directory Domain Controller	
Domain Name	Type the Active Directory domain name	
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 (Access	Policy>SSO Configurations)	
SSO Method	Kerberos	
Kerberos Realm	Type the Kerberos Realm. This must be uppercase, such as MYDOMAIN.COM	
KDC ¹	IP address of the Kerberos Key Distribution Center. If you leave this field blank, the system uses DNS to find the address of the KDC	
Account Name	The account name of the Active Directory user account to which logon rights have been delegated; this must begin with host/, for example, host/bigip_user_acct.mydomain.local	
Account Password	Type the associated password	
SPN Pattern ¹	Optional: Specify a custom SPN pattern to create the ticket request using the host name from the HTTP request ¹ .	
NTLM Machine Account (Ac	ccess Policy>Access Profiles>NTLM)	
Machine Account Name	The name of the account which will be joined to the Active Directory domain. This must be different than the account name specified in Kerberos SSO Configuration (such as bigip_machine_acct). Do not use spaces or special characters.	
Domain FQDN	Type the FQDN for Active Directory (such as mydomain.com)	
Admin User	Type the user name of a user with permissions to join a computer account to the Active Directory domain.	
Admin Password ¹	Type the associated password.	
NTLM Auth Configuration ²	(Access Policy> Access Profiles>NTLM)	
Name	Use following syntax: exch_ntlm_ <vs-name>, i.e. exch_ntlm_my_exchange_iapp_combined_https</vs-name>	
Machine Account Name	Select the NTLM Machine Account you created above	
Domain Controller FQDN List	Type the fully qualified name of your Active Directory domain controller and then click Add.	
11.4 and later only: Exchange	ge Profile ³ (Main tab>Access Policy>Secure Connectivity> Application Access> Microsoft Exchange)	
Parent Profile	/Common/exchange	
NTLM Configuration	Select the NTLM Auth configuration you created.	
In the left _i	pane of the box, click Autodiscover	
SSO Configuration	From the Autodiscover SSO Configuration list, select the Kerberos SSO Configuration you created above.	
In the left _i	pane of the box, click Exchange Web Service	
SSO Configuration	From the EWS SSO Configuration list, select the Kerberos SSO Configuration you created above.	
In the left _i	pane of the box, click Offline Address Book	
SSO Configuration	From the OAB SSO Configuration list, select the Kerberos SSO Configuration you created above.	
In the left pane of the box, click Outlook Anywhere		
Front End Authentication	ont End Authentication NTLM	
SSO Configuration	From the OAB SSO Configuration list, select the Kerberos SSO Configuration you created above.	
Access Profile (Main tab>Access Policy>Access Profiles)		
Microsoft Exchange	If you created the Exchange profile, select the profile you created from the list.	
SSO Configuration	Select name of Kerberos SSO configuration you created above	
Edit the Access Policy		
Edit the Access Profile you just or	reated using the Visual Policy Editor. Continue now with Editing the Access Policy on page 103	

By default, the SSO will attempt to use reverse DNS lookups of the pool member IP address to construct the Kerberos ticket request. If you do not wish to use DNS to find the host name to be used in the ticket request, you can specify a custom SPN pattern to create the ticket request using the host name from the HTTP request. The correct SPN pattern is: HTTP/%h@REALM.COM, where REALM.com is replaced with your fully-qualified Active Directory domain name. This configuration also requires that the DefaultAppPool, MSExchangeAutodiscoverAppPool, and MSExchangeServicesAppPool IIS application pools are configured to run under the user account specified for Kerberos Delegation, and that an SPN has been created for the hostname used to access Outlook Anywhere and Autodiscover 2 you must create this object in the same partition and folder location as the virtual server to which the Access Profile is applied. if you are manually reconfiguring the BIG-IP system from a previous iApp deployment, you will need to create this object from the tmsh command line. See the following procedure.
If using the Exchange profile in 11.4 and later, you must remove any _sys_APM irules from the virtual server.

F5 Deployment Guide

Creating the NTLM Auth Configuration from the TMSH command line

As noted in the preceding table, you must create the NTLM Auth Configuration object in the same partition and folder location as the virtual server to which the Access Profile is applied.

If you are manually reconfiguring the BIG-IP system from a previous iApp deployment, you need to create this object from the **tmsh** command line. This is only necessary if configuring the BIG-IP system from a previous iApp deployment.

To create the NTLM Auth Configuration from the command line

- 1. Open a command line session to the BIG-IP system
- 2. Type tmsh and then press Enter.
- 3. Type the command, using the following command syntax:

create apm ntlm ntlm-auth <iapp-name>.app/exch_ntlm_<virtual-server-name> app-service <iapp-name>.app dc-fqdnlist add { <domain-controller-fqdn> } machine-account-name <ntlm-machine-account-name>

For example, if the iApp is named **my_exchange_iapp**, the domain controller FQDN is **dc.mydomain.com**, the machine account is **bigip_machine_acct**, and the virtual server is named **my_exchange_iapp_combined_https**, the tmsh commands is:

create apm ntlm ntlm-auth my_exchange_iapp.app/exch_ntlm_my_exchange_iapp_combined_https app-service my_
exchange_iapp.app dc-fqdn-list add { dc.mydomain.com } machine-account-name bigip_machine_acct

Editing the Access Policy

The configuration in this section depends on whether you configured a separate virtual server for Outlook Anywhere, or configured a combined virtual server.

Editing the Access profile for Outlook Anywhere on a separate virtual server

Use the following procedure for configuring the Access Policy for a separate Outlook Anywhere virtual server.

To configure the Access Policy for Outlook Anywhere on a separate virtual server

- 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.
- 3. Click the + symbol between Start and Deny. A box opens with options for different actions.
 - a. Click the Client for MS Exchange option button, and then the Add Item button at the bottom.
 - b. Click the **Save** button.
- 4. On the Client for MS Exchange path, click the + symbol between Client for MS Exchange and Deny. A box opens with options for different actions.
 - a. Click the NTLM Auth Result Check option button, and then the Add Item button at the bottom.
 - b. Click the Save button.
- 5. On the Successful path between NTLM Auth Result Check 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**.
- 7. On the Fallback path between NTLM Auth Result Check and Deny, click the + symbol.
 - a. Click the **Logon Page** option button, and then the **Add Item** button at the bottom.
 - b. From the Split domain from full Username list, select Yes.
 - c. Configure the rest of the Logon Page properties as applicable, and then click **Save**.

- 8. On the Fallback path between Logon Page and Deny, click the + symbol.
 - 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.
 - c. Click Save.
- 9. 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.
 - Configure the Properties as applicable for your configuration. In our example, we leave the settings at the defaults. Click the Save button.
- 10. On the *fallback* path between **SSO Credential Mapping** and **Deny**, click the **Deny** box. Click the **Allow** option button, and then click **Save**.
- 11. 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. Your VPE should look like the following example.

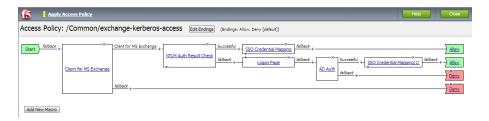


Figure 10: Example of the Access Policy in the VPE

This completes the Access Policy for the separate virtual server scenario. Continue with <u>Enabling the ECA Profile on Outlook</u> Anywhere Virtual Server on page 106.

Editing the Access profile for Outlook Anywhere on a combined virtual server

Use the following procedure for configuring the Access Policy if you configured Outlook Anywhere as a part of a combined virtual server.

To configure the Access Policy for Outlook Anywhere on a combined virtual server

- 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.
- Click the + symbol between Start and Deny. A box opens with options for different actions.
 - a. Click the Client for MS Exchange option button, and then the Add Item button at the bottom.
 - b. Click the **Save** button.
- 4. On the Client for MS Exchange path, click the + symbol between Client for MS Exchange and Deny. A box opens with options for different actions.
 - a. Click the NTLM Auth Result Check option button, and then the Add Item button at the bottom.
 - b. Click the Save button.
- 5. On the Successful path between NTLM Auth Result Check 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**.

- 7. On the Fallback path between NTLM Auth Result Check and Deny, click the + symbol.
 - a. Click the Logon Page option button, and then the Add Item button at the bottom.
 - b. From the Split domain from full Username list, select Yes.
 - c. Configure the rest of the Logon Page properties as applicable, and then click Save.
- 8. On the Fallback path between Logon Page and Deny, click the + symbol.
 - 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.
 - c. Click Save.
- 9. 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.
- 10. On the *fallback* path between **SSO Credential Mapping** and **Deny**, click the **Deny** box. Click the **Allow** option button, and then click **Save**.
- 11. On the Fallback path between Client for MS Exchange (the first box of the VPE) and Deny, click the + symbol.
 - a. Click the **Logon Page** option button, and then the **Add Item** button at the bottom.
 - b. From the Split domain from full Username list, select Yes.
 - c. Configure the rest of the Logon Page properties as applicable, and then click Save.
- 12. On the bottom Fallback path between the new Logon Page and Deny, click the + symbol.
 - 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.
 - c. Click Save.
- 13. 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.
- 14. On the *fallback* path between **SSO Credential Mapping** and **Deny**, click the **Deny** box. Click the **Allow** option button, and then click **Save**.
- 15. 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. Your VPE should look like the following example.



Figure 11: Example of the Access Policy in the VPE

This completes the Access Policy for the combined virtual server.

Enabling the ECA Profile on Outlook Anywhere Virtual Server

The next task is to enable the ECA profile on the Outlook Anywhere virtual server. This profile allows the APM to manage NTLM authentication for Outlook Anywhere clients. In BIG-IP version 11.3, you must attach the ECA profile via the **tmsh** command line. Do NOT enable this profile if using BIG-IP version 11.4 and the Exchange Profile.

To attach the ECA profile to the virtual server from the command line

- 1. Open a command line session to the BIG-IP system
- 2. Type tmsh
- 3. Type the command, using the following command syntax:

```
modify ltm virtual <iapp-name>.app/<virtual-server-name> profiles add { eca }
For example:
```

modify ltm virtual my_exchange_iapp.app/my_exchange_iapp_combined_https profiles add { eca }

(i) Important You must have enabled the ECA profile on the Outlook Anywhere virtual server as described above before applying the system iRule in the next step. Do NOT enable this profile or attach the system iRule if using BIG-IP version 11.4 and the Exchange Profile.

Applying the System iRule to Outlook Anywhere virtual server if using a BIG-IP version prior to 11.4

Before attempting a connection via BIG-IP APM with Outlook Anywhere, you must apply the system iRule that manages NTLM authentication to either the separate Outlook Anywhere virtual server, or the combined virtual server.

To apply the system iRule to the virtual server

- 1. On the Main tab, expand Local Traffic and then click Virtual Servers.
- 2. Click the name of either the combined virtual server or the separate Outlook Anywhere virtual server.
- 3. Click the **Resources** tab.
- 4. In the iRules section, click the **Manage** button.
- 5. From the Available list, select _sys_APM_ExchangeSupport_OA_NtlmAuth, and then click the Add (<<) button.
- 6. If necessary, use the Up and Down buttons to ensure the iRules are in the following order when deployed on a single, combined virtual server:
 - OWA Append iRule (for combined virtual only)
 - _sys_APM_ExchangeSupport_OA_NtlmAuth
 - Select SSO iRule
 - Combined Virtual Server Persist iRule
- 7. Click Finished.

Setting the Default Pool on a combined virtual server

If you have configured the BIG-IP system use a single, combined virtual server for Exchange, the final task is to set the BIG-IP LTM pool for Outlook Anywhere as the default pool for the virtual server.

To set the default pool on the combined virtual server

- 1. On the Main tab, expand Local Traffic and then click Virtual Servers.
- 2. Click the HTTPS virtual server (port 443) virtual server.
- 3. Click the **Resources** tab.
- 4. In the Load Balancing section, from the **Default Pool** list, select the Outlook Anywhere pool.
- 5. Click Update.

This completes the configuration for NTLM and Outlook Anywhere.

Access Policy example when using both EAC restricted access and NTLM for Outlook Anywhere

Using both EAC restricted access and NTLM for Outlook Anywhere in a single Access Policy is an acceptable configuration, although the step by step procedure is outside the scope of this document (use the iApp for this scenario if you need the walkthrough). The following screenshot shows what the VPE should look like with both EAC restricted access and NTLM for Outlook Anywhere.

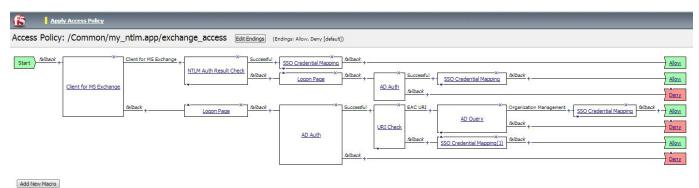


Figure 12: Example of the Access Policy in the VPE

Manually configuring the BIG-IP Advanced Firewall Module to secure your Exchange deployment

This section describes how to manually configure BIG-IP AFM, F5's Network Firewall module, to secure your Exchange deployment. BIG-IP AFM is particularly useful if you want to only allow access from specific clients or networks. Because this configuration can be complex, we recommend using the iApp template in version 11.6 and later to configure BIG-IP AFM.

Network Firewall settings

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

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

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

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

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

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

3. Creating a firewall rule to block all other traffic

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

- a. Click Security > Network Firewall > Policies.
- b. Click the name of the policy you created in step 1.
- c. In the Rule section (below the General Properties section), click the Add button.
- d. Leave the **Type** list set to **Rule**.
- e. Leave the Order list, select Last.
- f. In the Name field, type a unique name, for example Exchange-traffic-Prohibited.
- g. Ensure the State list is set to Enabled.
- h. From the **Protocol** list, select **TCP**. Leave the box to the right of TCP set to **6**.
- i. In the **Source** section, leave all the lists set to **Any**
- j. From the **Action** list, select either **Drop** (to silently discard incoming connections) or **Reject** (to send a Destination Unreachable message to the sender).
- k. If you configured a logging profile as described in <u>Optional: Configuring the BIG-IP system to log network firewall events on page 110</u>, from the **Logging** list, select **Enabled**.
- I. Click **Finished**. You return to the Policy Properties page.
- m. On the Policy Properties page, in the Rules section, ensure the rule with the Action of Accept comes before the Drop or Reject rule you just created. If it does not, use the **Reorder** button and drag the rules into the correct order.

4. Apply Your Firewall Policy to your Virtual Server

- a. Click Security > Network Firewall > Active Rules.
- b. In the Rule section (below the General Properties section), click the Add button.
- c. From the Context list, select Virtual Server, and then select the virtual server you created for your Exchange traffic.
- d. From the Type list, select Policy, and then select the firewall policy you created.
- e. From the **Policy Type** list, select **Enforced**.
- f. Click **Finished**.

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

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

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

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

To assign the IP intelligence policy to the Exchange virtual server

- 1. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.
- 2. Click the name of your Exchange virtual server.

- 3. From the Security menu, choose Policies.
- 4. Next to IP Intelligence, select Enabled, then select the IP intelligence policy to apply to traffic on the virtual server.
- 5. Click **Update**. The list screen and the updated item are displayed. The IP Intelligence policy is applied to traffic on the virtual server.

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

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

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

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

Creating the logging profile using the iApp template

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

To configure the logging profile iApp

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

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

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

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

Creating logging profile manually

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

To manually configure a logging profile

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

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

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

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

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

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

5. Create a log publisher:

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

6. Create the logging profile to tie everything together.

If you chose to log allowed connections, include the green text (as in step 2 substep I in <u>To configure the BIG-IP AFM to allow connections from a single trusted network on page 108</u>).

If you set the rule to drop incoming connections, include the text in blue.

If you chose to log IP intelligence events, include the text in red to add the parameter that sets the log publisher.

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

Assigning the logging profile to the virtual server

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

To assign the logging profile to the Exchange virtual server

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

Appendix D: Technical Notes

The following contains additional information that may be helpful when configuring the BIG-IP system for Exchange Server 2010 and 2013.

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 Client Access server to fully start all services and be ready to handle a full load of traffic, but that time is highly dependent on local conditions. You may want to adjust the time period up or down in your environment based on your server capacity and load.

Subject Alternative Name (SAN) SSL Certificates

This template currently only supports the use of a single DNS name and corresponding certificate and key for all services, or multiple DNS names using a SAN-enabled certificate and key or wildcard certificate.

An SSL certificate that supports the Subject Alternative Name (SAN) extension allows more than one valid FQDN per certificate, without having to resort to a "wildcard" certificate for a domain. When used in conjunction with Exchange Server, 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 here:

http://technet.microsoft.com/en-us/library/aa995942%28EXCHG.140%29.aspx

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

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

The BIG-IP system supports using a wildcard certificate to secure Exchange CAS deployments using multiple FQDNs. However, for increased security, F5 recommends using SAN certificate(s) where possible. Additionally, some older mobile devices are incompatible with wildcard certificates. Consult your issuing Certificate Authority for compatibility information.

Maximum number of concurrent users: SNAT Pool guidance

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

The BIG-IP system can create roughly 64,000 connections per SNAT address (ephemeral or source ports used by connections from the BIG-IP range from 1024 to 65,535, or an absolute maximum 64,511 effective concurrent connections). Each user connected to a Client Access server can have about 10 concurrent connections (for example, if a user has Outlook on a PC, a mobile phone, and Lync running simultaneously). Therefore, you would need a SNAT address for each 6,000 concurrent users you expect. For example, if you have 12,000 users, you need two SNAT pool IP addresses; if you have 15,000 users, you need three addresses. The IP address(es) you specify must not be self IP addresses on this BIG-IP system.

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 BIG-IP APM	
Connect using SSL only	Selected	Selected	
On fast networks, connect using HTTP first, then connect using TCP/IP	Not selected	Selected	
On slow networks, connect using HTTP first, then connect using TCP/IP	Selected	Selected	
Proxy authentication settings	NTLM	Basic	

Creating a new Client Access Array

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

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

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

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

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

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

For complete documentation from Microsoft, see

http://technet.microsoft.com/en-us/library/ee332317.aspx

Note on creating advanced monitors manually

If you choose advanced monitors, the BIG-IP system performs logins to most of the Client Access services (all except RPC/MAPI in Exchange 2010 and Forms-based Outlook Web App) and checks for valid content in the response. Because these monitors attempt to access a specific mailbox, they more accurately determine the actual health of CAS services. However, account maintenance and Mailbox status must become a part of your monitoring strategy.

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

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

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

You should create accounts (and associated mailboxes) for monitoring that are not accessed by actual users and that do not have privileged access anywhere else in your network. Because you have to store the user name and password in plain text on your BIG-IP devices, make sure the credentials are not used elsewhere in your organization for anything other than monitoring. We strongly recommend creating a mailbox account(s) specifically for use in the monitor(s).

Appendix E: Active Directory and Exchange Server configuration for NTLM

If you plan on configuring your BIG-IP system version for NTLM authentication as described in Optional: Configuring the APM for Outlook Anywhere with NTLM Authentication - BIG-IP v11.3 or later only on page 102, you must first perform the following tasks on your Active Directory and Exchange servers.

Most of the following guidance is performed using Microsoft PowerShell. You must have access to perform PowerShell commands.

Note: This section provides guidance only; for specific instructions, consult the appropriate documentation. F5 cannot be responsible for improper configuration of Active Directory or Microsoft devices.

Create a Delegation Account

You must create a user account for the BIG-IP system to use to perform Kerberos authentication. The user logon name must begin with host/ and the account should be a member of the Domain Users security group. If you are deploying NTLM for Outlook Anywhere and Exchange 2010 without using DNS lookups (BIG-IP APM/LTM without DNS lookups on page 117), the delegation account must also be a member of the Exchange Trusted Subsystem, Exchange Windows Permissions, and IIS_USRS Active Directory security groups. If you are deploying Exchange 2013, the Alternate Service Account does not need to be a member of these additional security groups.

Run the following PowerShell commands on an Active Directory Domain Controller on a single line, replacing the text in red with the proper information for your environment. You will be prompted to enter a new password for the delegation account.

New-ADUser -Name "APM Delegation Account" -UserPrincipalName host/account-username.example.com@example.com -SamAccountName "account-username" -PasswordNeverExpires \$true -Enabled \$true -AccountPassword (Read-Host -AsSecureString "Account Password")

Next, add the user account to the required Active Directory groups

Add-ADGroupMember "Exchange Trusted Subsystem" account-username Add-ADGroupMember "Exchange Windows Permissions" account-username Add-ADGroupMember "IIS_IUSRS" account-username

Configure the servicePrincipalName

The next task is to modify the servicePrincipalName attribute of the Delegation Account. The servicePrincipalName value should match the user logon name of the delegation account. Replace the domain in the example with your domain

Set-AdUser -Identity account-username -ServicePrincipalNames @{Add="host/account-username.example.com"}

Enabling Delegation for the account

After configuring the servicePrincipalName attribute, find the account in Active Directory Users and Computers. After configuring the servicePrincipalName attribute, the Delegation tab appears under the properties of the user account. Select Trust the user for delegation to specified services only, and then select Use any authentication method. Click Add to add a service for which this account can authenticate, and then add the HTTP service type for each Client Access Server.

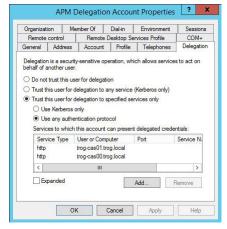


Figure 13: APM Delegation Account properties

Configure Outlook Anywhere for NTLM Client Authentication

Run the following commands on one or more Exchange Client Access Servers to configure Outlook Anywhere for NTLM. You may run this command on any Client Access Server.

Get-OutlookAnywhere | Set-OutlookAnywhere -ExternalClientAuthenticationMethod NTLM -InternalClientAuthenticationMethod NTLM

Enabling Kerberos Authentication for RPC IIS Virtual Directory

Enable the Negotiate authentication provider on the RPC virtual directory using the following PowerShell command:

Get-OutlookAnywhere | Set-OutlookAnywhere -IISAuthenticationMethods Negotiate,NTLM

DNS Reverse Lookups

If the Outlook Anywhere IIS Application Pool is running under the LocalSystem or ApplicationPoolIdentity account, you must ensure that APM can successfully perform reverse DNS lookups against the IP address of the Outlook Anywhere pool member(s). These DNS lookups must return the host name of the Exchange CAS server (APM+LTM scenario):

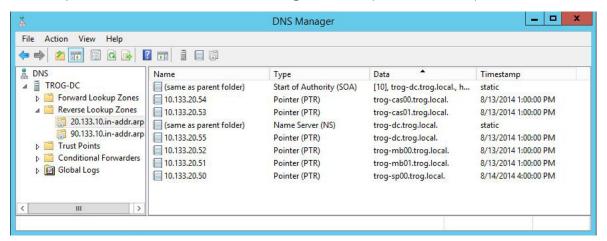


Figure 14: DNS Manager

BIG-IP APM/LTM without DNS lookups

If you have deployed BIG-IP APM to forward Outlook Anywhere traffic to a virtual server on an internal BIG-IP LTM, or you are deploying on a BIG-IP system running both LTM and APM and would like to eliminate the need for reverse DNS lookups, you must perform the following configuration steps in Active Directory and from the IIS Management Console on the Client Access Servers.

The first task in this section is to create a Service Principal Name for the Outlook Anywhere FQDN to allow authentication by the delegation user account. Replace the text in red with the appropriate values in your implementation.

setspn -S <SPN> <ACCOUNT>
setspn -S http/mail.example.com EXAMPLE\account-username
setspn -S http/autodiscover.example.com EXAMPLE\account-username

Perform this step for every host name that you will be accessing using NTLM client authentication, which includes Autodiscover by default.

Based on reverse DNS lookups or the SPN pattern specified in the Kerberos SSO configuration on page 98, APM will construct a Kerberos ticket request to the Active Directory domain controller for the SPN HTTP/mail.example.com. You must allow Kerberos constrained delegation for HTTP/mail.example.com via the Delegation tab within the properties of the previously created user account.

Also, you must ensure that the previously created delegation account is allowed to log on for all of the SPNs you just created (see *Enabling Delegation for the account on page 115*).

Finally, you must change the Application Pool Identity for the Application Pool used by Outlook Anywhere, Autodiscover, and Exchange Web Services to use the delegation user account you created, or configure an Alternate Service Account for each Client Access Server.

Setting the IIS Application Pool Identity or Alternate Service Account

Use the following PowerShell commands to set the IIS Application Pool Identity for Exchange Server 2010, or the Alternate Service Account for Exchange Server 2013. Use the commands for the version of Microsoft Exchange Server you are using; 2010 or 2013.

Setting the IIS Application Pool Identity for Exchange Server 2010

You must run the following PowerShell commands on all the Exchange 2010 Client Access Servers in your implementation.

Import the IIS PowerShell module:

Import-Module WebAdministration

Check the App Pool configured for Outlook Anywhere

Get-WebApplication "rpc"

Set the Application Identity for the App Pool returned by the previous command

cd IIS:\AppPools

Set-ItemProperty "DefaultAppPool" -Name processodel.identityType -Value 3

Set-ItemProperty "DefaultAppPool" -Name processodel.username -Value "EXAMPLE\account-username"

Set-ItemProperty "DefaultAppPool" -Name processodel.password -Value "<APM delegation account password>"

Repeat these commands for the Autodiscover and Exchange Web Services virtual directories.

Configuring Alternate Service Account for Exchange 2013 Client Access Servers

The commands for configuring the Alternate Service Account for Exchange 2013 depends on whether your Client Access and Mailbox roles are located on the same servers or not.

If your Client Access Servers are **separate** from your Mailbox servers

You must run the following PowerShell commands on all Client Access Servers in your implementation. In the following examples, <CAS_FQDN> is the fully-qualified name of the Client Access Server, and <CAS> is the short name of the Client Access Server

```
$cred = Get-Credential
```

\$session = CreateOrGetExchangeSession <CAS_FQDN> \$null \$true \$false <CAS_FQDN>
Invoke-Command -Session \$session -Args (\$cred) -ScriptBlock {param(\$AsaCreds) Set-ClientAccessServer <CAS>
-AlternateServiceAccountCredential \$AsaCreds}

If your Client Access and Mailbox roles are on the same servers

If your Client Access and Mailbox roles are on the same servers, you must run these commands on all combined Client Access/Mailbox Servers in your deployment. In the following example, <CAS_MBX> is the short name of the co-located Client Access/Mailbox Server.

Set-ClientAccessServer <CAS_MBX> -AlternateServiceAccountCredential (Get-Credential)

Finally, verify ASA. You may run this command on any Client Access Server in the implementation.

Get-ClientAccessServer -IncludeAlternateServiceAccountCredentialStatus | Format-List Name,Alt*

Troubleshooting NTLM Authentication

You can increase the logging level for Access Policy Manager to assist in troubleshooting issues with NTLM client authentication. Click **System > Logs > Configuration > Options**. Under **Access Policy Logging**, select **Debug** log level for either the Access Policy, SSO, or both. The debug setting causes BIG-IP to log all APM-related messages to this file: /var/log/apm.

These logs can be useful in diagnosing problems with NTLM auth/Kerberos SSO functionality.

If you have followed these steps and are receiving Kerberos errors in the APM log, you can clear any previously cached Kerberos tickets by restarting the websso service on the APM BIG-IP system:

[root@ms-ve-v11-x2010-EDGE:Active:Standalone] config # bigstart restart websso

Document Revision History

Version	Description	Date
	New deployment guide for the fully supported f5.microsoft_exchange_2010_2013_cas.v1.5.0 iApp template. This iApp contains the following major enhancements:	
	- Added support for BIG-IP Advanced Firewall Manager (AFM)	
1.0	- Added the ability for password change from APM/OWA logon page	04-09-2015
	- Removed the option to choose a pre-existing direct AAA server object	
	- Modified the troubleshooting entry, <u>Lync clients cannot connect or receive authentication prompts when accessing Microsoft Exchange Autodiscover and EWS through F5 APM on page 62</u> with an updated iRule.	
1.1	- Added a new section, <u>Modifying the iApp configuration on page 52</u> with required changes if you configured the iApp to use different IP addresses for different services, and are deploying ActiveSync and/or Outlook Anywhere.	04-22-2015
1.2	- Added a new troubleshooting entry: You may experience deployment errors when the NTLM Machine Account name contains spaces or special characters on page 64. Added the same guidance to the manual configuration.	05-01-2015
1.3	 - Added a new troubleshooting entry: <u>The Direct File Access setting for public computers is not honored on page 64</u>. - Added support for all CUs in the Product Version table. 	08-07-2015

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

F5 Networks, Inc. Corporate Headquarters info@f5.com F5 Networks Asia-Pacific apacinfo@f5.com F5 Networks Ltd. Europe/Middle-East/Africa emeainfo@f5.com F5 Networks Japan K.K. f5j-info@f5.com

