



Deploying the BIG-IP Edge Gateway for Layered Security and Acceleration Services

Table of Contents

Using the BIG-IP Edge Gateway for layered security and acceleration services	
Prerequisites and configuration notes	1
Product versions and revision history	1
Configuration example	2
Configuration flow	2
Performing the initial configuration tasks	3
Configuring access and acceleration for mobile users	
Configuring authenticated remote network access	6
Creating an AAA resource	8
Creating a Connectivity Profile	10
Creating a Webtop	10
Creating the Rewrite profile	11
Creating an Access Profile	11
Editing the Access Profile with the Visual Policy Editor	12
Configuring network entry points	14
Creating the profiles	14
Creating the virtual servers	16
Configuring site-to-site WAN optimization for application access across data centers	19
Prerequisites and configuration notes	19
Configuring the WAN Optimization module	19
Creating the iSession profile	19
Configuring BIG-IP WOM using the Quick Start template	20
Configuring authenticated access to accelerated web applications	23
Creating an AAA server	23
Creating the SSO configuration	23
Creating an Access Profile	24
Editing the Access Profile with the Visual Policy Editor	25
Creating the profiles	26
Creating the virtual server	27
Configuring the WebAccelerator	29

Using the BIG-IP Edge Gateway for layered security and acceleration services

Welcome to the accelerated remote access deployment guide. This guide shows how to configure the BIG-IP Edge Gateway to accelerate secure access to application resources. Using the configuration in this guide, the BIG-IP Edge Gateway performs client security checks, optimizes remote access connections and accelerates application access.

With the BIG-IP Edge Gateway, it is now possible to layer remote access and acceleration for network streams and web application objects. Edge Gateway enables you to achieve the fastest access times, using the least bandwidth and traversing the most expeditious path between application clients and servers.

This deployment places the BIG-IP Edge Gateway in an existing DMZ and connects networks of like security levels. The BIG-IP Edge Gateway will provide authentication, authorization, auditing, prelogon inspection services, and accelerate access to remote users of applications presented in a DMZ.

For more information on the BIG-IP Edge Gateway, see <http://www.f5.com/products/big-ip/solution-modules/edge-gateway.html>

Prerequisites and configuration notes

The following are prerequisites for this solution.

- ◆ The management port is configured and the web based configuration interface is accessible
- ◆ For WebAccelerator, we recommend you review the applications in your environment to determine if there is an existing WA policy and choose the best candidates for acceleration.
- ◆ For WAN Optimization, determine all the protocols used by your application that will need to be optimized (FTP, HTTP, CIFS, MAPI, etc).
- ◆ For Access Policy Manager you need to know where authentication and authorization services (RADIUS, Active Directory, LDAP etc) are on your network, as well as map your organization's security policy to prelogon checks.

Product versions and revision history

Product and versions tested for this deployment guide:

Product Tested	Version Tested
BIG-IP Edge Gateway	v10.2

Document Version	Description
1.0	New guide

Configuration example

The following diagram shows a logical configuration example of our deployment.

Configuration flow

The following chart for configuring remote access using an Access Policy (read from the bottom to the top) illustrates the setup of Network Access within Edge Gateway. The information about Web Application is included for reference but is not part of the setup for Network Access.

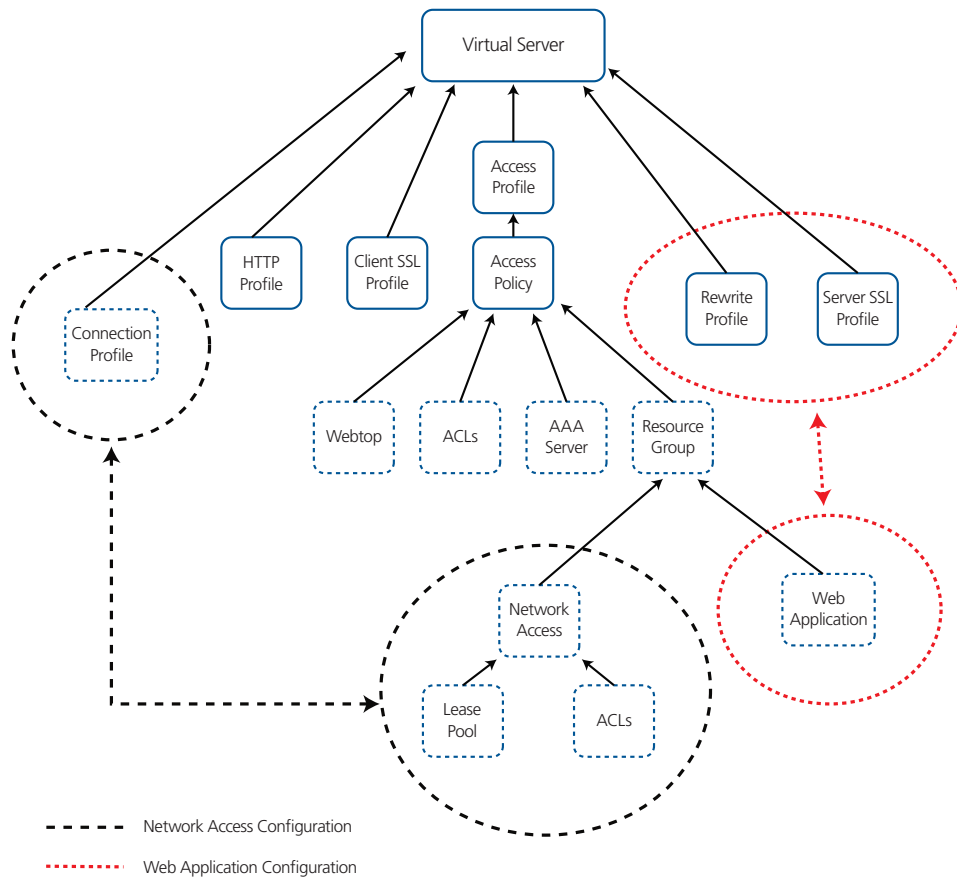


Figure 1 Configuration flow

Performing the initial configuration tasks

The first task is to perform the initial configuration tasks. Here we configure the BIG-IP system with required VLAN and Self IP address information. Complete these procedures only if you do not already configured these objects on the BIG-IP LTM.

Creating a VLAN

First we create a VLAN on the BIG-IP LTM system.

To create a VLAN

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

Creating self IP addresses

Self IP addresses are the IP addresses owned by the BIG-IP LTM system that you use to access the VLAN. The next step in this configuration is to create a self IP address that is used in the local end point BIG-IP configuration.

To create a self IP address

1. On the Main tab, expand **Network**, and then click **Self IPs**.
The Self IP screen opens.
2. Click the **Create** button.
The new Self IP screen opens.
3. In the **IP Address** box, type a static IP address that resides on the VLAN you created in the preceding procedure.
4. In the **Netmask** box, type the corresponding subnet mask.
5. From the **VLAN** list, select the VLAN you created in *Creating a VLAN*. In our example, we select **edge-vlan**.

6. From the **Port Lockdown** list, select **Allow Default**.
7. Click the **Repeat** button. Repeat steps 1-6 with the following exception:
 - In Step 6, from the **Port Lockdown** list, select **Allow None**.
8. Click **Finished**.
9. Repeat this entire procedure on the remote endpoint BIG-IP system.

Configuring Routes

The next task is to create the Routes on the BIG-IP system. Each Edge Gateway needs to be able to route to the other Edge Gateway, as well as have a route for the remote network where application services reside. We assume that the firewall is the default route (.1).

Additionally, each Edge Gateway should be able to route to application and AAA servers via the local internal firewall (.254).

For WAN optimization to occur properly between sites, each Edge Gateway will also need a route to the remote networks via the appropriate next hop gateway. In our example, we add a route for 10.10.2.0/24 via 192.168.2.1 in Datacenter 1.

To create the route

1. On the Main tab, expand **Network**, and then click **Routes**.
2. Click the Create button.
3. From the **Type** list, select **Route**.
4. In the **Destination** box, type the IP network address of the remote network you wish to reach.
5. In the **Netmask** box, type the associated Netmask.
6. From the **Resource** list, make sure **Use Gateway** is selected.
7. From the Gateway Address list, select **IP Address**, and then type the IP address of the next hop gateway used to access this network (see Figure 2).
8. Click **Repeat**. Repeat this procedure three times to create routes for the following:
 - The remote network
 - The default gateway
 - The internal network
9. Click **Finished**.

Network » Routes » New Route...

Properties

Type	Route
Destination	10.10.2.0
Netmask	255.255.255.0
Resource	Use Gateway...
Gateway Address	IP Address 192.168.2.1

Cancel Repeat Finished

Figure 2 New Route configuration

Configuring access and acceleration for mobile users

To configure remote access and acceleration, a Device Wizard is included in the product that assists in the setup of Network Access. In this guide, we describe the steps to complete the configuration manually. In this section we cover three scenarios:

1. *Configuring authenticated remote network access*, on page 6
2. *Configuring site-to-site WAN optimization for application access across data centers*, on page 19
3. *Configuring authenticated access to accelerated web applications*, on page 23

Configuring authenticated remote network access

The authenticated remote network access scenario results in a configuration where users are able to authenticate to the BIG-IP Edge Gateway and access private network resources via a private network tunnel between the client and the BIG-IP Edge Gateway.

Use the following procedures to configure authenticated remote network access.

To configure remote access

1. On the **Main** tab, expand **Access Policy**, and then click **Network Access**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this Network Access Profile. In our example, we type **London_Remote_Access**. You can optionally type a description.
4. In the General Settings section, next to **Lease Pool**, click the Add (+) button. The Lease Pool is the pool of IP Addresses that clients receive when they connect to the VPN.
 - a) In the **Name** box, type a name for the Lease pool. In our example, we type **London_Lease_Pool**.
 - b) Click the **IP Address Range** button.
 - c) In the **Start IP Address** and **End IP Address** boxes, type the appropriate IP addresses. In our example, we allow addresses from **10.0.1.1** to **10.0.1.255**.
 - d) Click the **Add** button.

5. Click the **Finished** button. You return to the Network Access list.

Access Policy >> Network Access : Lease Pools >> New Lease Pool...

General Properties

Name: London_Lease_Pool

Configuration

Type: IP Address IP Address Range

Start IP Address: 10.0.1.1

End IP Address: 10.0.1.255

Add

Member List

10.0.1.1 - 10.0.1.255

Edit Delete

Cancel Finished

Figure 3 Configuring the Lease Pool

6. If necessary, from the **Lease Pool** list, select the lease pool you just created. In our example, we select **London_Lease_Pool**.
7. To enable optimization on the client tunnel itself, from the **Compression** list, select **GZIP Compression**. This allows both the web browser client and the BIG-IP Edge client to take advantage of compression between the client and the remote access server.

Note: If DTLS is configured (UDP based communication between client and Remote Access Server) GZIP compression is automatically disabled on a per-user-session basis. If D/TLS and Gzip compression are both enabled and the client is unable to connect with D/TLS, Gzip compression will be used on a per-user-session basis.

Access Policy >> Network Access : Network Access List >> New Resource...

General Properties

Name: London_Remote_Access

Description:

General Settings: Advanced

Lease Pool: London_Lease_Pool

Compression: GZIP Compression

SNAT Pool: Auto Map

Session Update Threshold: 0

Session Update Window: 0

Figure 4 Configuring Network Access

8. From the **Client Settings** list, select **Advanced**.
9. In the Traffic Options section, you can choose to Force all traffic through the tunnel, or use split tunneling. With Split Tunneling enabled, the administrator needs to indicate which subnets should be routed through the VPN tunnel. If Split tunneling is not allowed, all traffic will go through the tunnel.
 - a) If you want all traffic to go through the tunnel, click **Force all traffic through tunnel**, and continue with Step 10.
 - b) If you want to use split tunneling, click **Use split tunneling** for traffic. The split tunneling options appear.
 - In the LAN Address Space section, type the IP address and Mask of the LAN Address space that should go through the tunnel. In our example we indicate that 192.168.0.0/16 is all LAN space. To allow clients to access applications in Datacenter 2, be sure that you add another entry for the DMZ subnet in the remote datacenter. We configure WOM to accelerate this access later on in this guide.
 - In the DNS Address Space section, type the DNS name(s) that are used in the target LAN.
 - In the Exclude Address Space section, type the IP address and Mask of any address space that should be excluded. For example, if a portion of 192.168.0.0/16 should be excluded, it can be entered here. In our example, we indicate that 192.168.10.0/24 is excluded.
10. The remaining options are also administrative, configure the settings as applicable to your configuration. In our testing and architecture we generally recommend the following settings:
 - a) In the Client Side Security section, we select **Prohibit routing table changes during Network Access Connection**.
 - b) In the Reconnect To Domain section, we select **Synchronize with Active Directory policies on connection establishment**.
 - c) In the DTLS section, check the box to enable DTLS. We recommend using DTLS protocol for optimum performance.

*Note: DTLS uses UDP port 4433 by default. Arrange to open this port on firewalls as needed.
For DTLS, a UDP Virtual Server is required (described in Creating the virtual servers, on page 16).*
11. Click **Finished**.

Creating an AAA resource

The Edge Gateway does not have a built-in authentication store therefore an authentication source must be specified. You can configure AAA in two ways:

-
- ◆ *Central AAA*
With central AAA, all of the Edge Gateways in your configuration point at a single authentication source.
 - ◆ *Local AAA*
With local AAA, each Edge Gateway connects to a local AAA resource, where the resources are replicated from site to site.

To create an AAA server

1. On the Main tab, expand **Access Policy**, and then click **AAA servers**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this profile. In our example, we type **Seattle_LDAP_server** because we traverse the wide-area-network back to Seattle to perform authentication.
4. From the **Type** list, select the appropriate authentication method. For this example, we select **LDAP**.
5. In the Configuration section, type the appropriate information relevant to your authentication method. In our LDAP example, we provide the Host name for the LDAP server, the Admin DN, and the Admin Password. We leave the timeout at default.
6. Click **Finished**.

Access Policy >> AAA Servers >> New Server...	
General Properties	
Name	Seattle_LDAP_server
Type	LDAP
Configuration	
Host	192.0.2.127
Service Port	389 LDAP
Admin DN	admin
Admin Password
Verify Admin Password
Timeout	15 seconds
Cancel Finished	

Figure 5 New AAA server configuration

Creating a Connectivity Profile

The next task is to create a connectivity profile.

To create a connectivity profile

1. On the Main tab, expand **Access Policy**, and then click **Connectivity Profile**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this profile. In our example, we type **London_Connectivity**.
4. Configure the rest of the options as applicable to your configuration. In our example, we leave all settings at the default.
5. Click **Finished**.

Creating a Webtop

In the BIG-IP Edge, a Network Webtop is used to deliver the BIG-IP Edge client components to the user's web browser session.

To create a Webtop

1. On the Main tab, expand **Access Policy**, and then click **Webtops**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this webtop. In our example, we type **London_Webtop**.
4. From the **Type** list, select **Network Access**.
5. If you want the browser window to be minimized to the system tray for Windows hosts, check the **Enabled** box.
6. Click **Finished**.

Access Policy >> Webtops >> New Webtop...	
General Properties	
Name	London_Webtop
Type	Network Access
Configuration	
Minimize To Tray	<input checked="" type="checkbox"/> Enabled
Cancel Repeat Finished	

Figure 6 New Webtop configuration

Creating the Rewrite profile

If you are providing access to applications separately from a network tunnel you will need to create a Rewrite Profile. The Rewrite profile allows BIG-IP Edge Gateways APM component to act as a reverse proxy.

To create the Rewrite profile

1. On the Main tab, expand **Access Policy**, and then click **Rewrite Profiles**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this profile. In our example, we type **edge-Rewrite**.
4. Leave the **Client Caching Type** list at the default (**CSS and JavaScript**).
5. Click the **Finished** button.

Creating an Access Profile

The Access Profile ties together the other configuration elements to create a Network Connection VPN Tunnel. The Access Profile is where the Visual Policy Editor (VPE) is located, which allows for workflow design to ensure appropriate security levels for each client access.

To create an Access Profile

1. On the Main tab, expand **Access Policy**, and then click **Access Profiles**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this profile. In our example, we type **London_Access_Policy**.
4. In the Settings section, configure the options as applicable for your configuration. In our example, we leave all of the settings at their defaults. Note that depending on licensing, the number of concurrent users may be limited. The other timeouts are administrative choices.
5. In the Configuration section, configure the settings as applicable to your environment. In our example, we accept all of the defaults. We are not using Single-Sign-On configurations or specific Logout URIs. However, we do leave **Secure Cookie** checked.
6. In the Language Settings section, if you are configuring the Edge Gateway in a language other than English, configure as applicable for your language. In our example, we accept English as the default language.
7. Click **Finished**.

Editing the Access Profile with the Visual Policy Editor

The next task is to open the London Access Policy and edit the Access Policy using the Visual Policy Editor (VPE). The VPE is a powerful visual scripting language that offers virtually unlimited options in configuring an Access Policy. For detailed information on the VPE please see the product documentation.

In the following procedure, we configure a policy using the Visual Policy Editor. However, Device Wizards provide an easy way to create more interesting policies, including ones that check for virus software and other prerequisites before allowing a user to logon. In this guide, it is our goal to get you oriented with the concepts of the Visual Policy Editor. In our example, we create a Login Page, LDAP auth, and assign the resources allowed.

To edit the Access Profile

1. On the Main tab, expand **Access Policy**, and then click **Access Profiles**.
2. Locate the Access Profile you just created, and in the Access Policy column, click **Edit**. The Visual Policy Editor opens.
3. Click the + symbol between **Start** and **Deny**. A box opens with options for different actions.
4. Click the **Logon Page** option button, and then click the **Add Item** button at the bottom of the box.
5. Configure the Properties as applicable for your configuration. In our example, we leave the settings at the defaults. Click the **Save** button.
6. Click the + symbol between **Logon Page** and **Deny**.
7. In the Authentication section, click the **LDAP Auth** option button, and then click the **Add Item** button.
8. From the **Server** list, select the AAA Source you created in *Creating an AAA resource*, on page 8 (see Figure 7).
9. Add **SearchDN** and **SearchFilter** items as applicable for your configuration.
10. Click the **Save** button. You now see two paths, **Successful** and **Fall Back**.

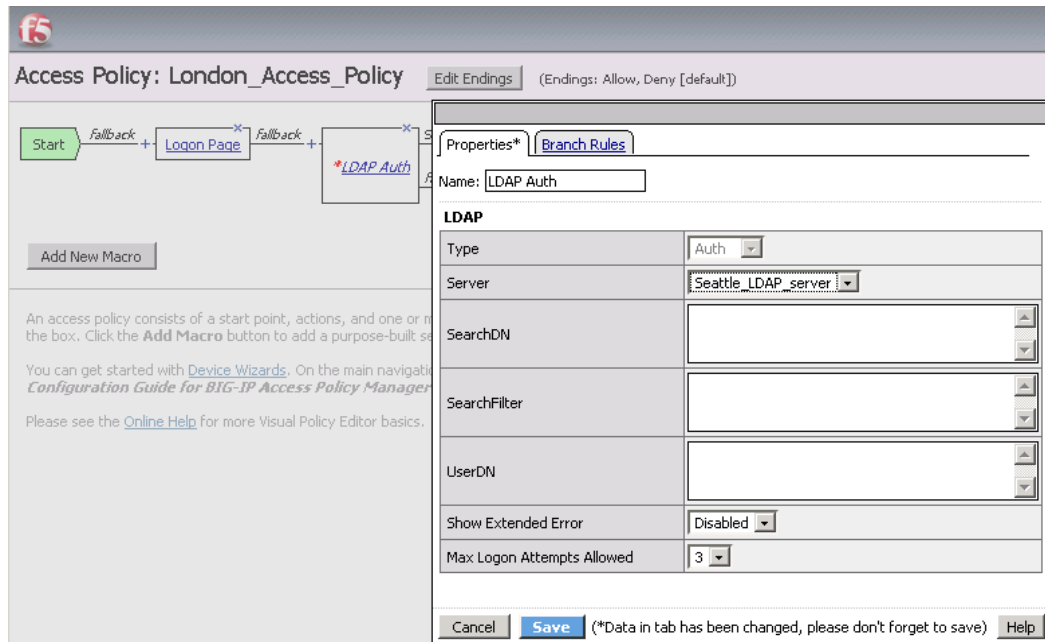


Figure 7 LDAP Authentication box on the Visual Policy Editor

11. Click the **Deny** box from the path leading from Successful. The Select Ending box opens.
12. Click the **Allow** button, and then click **Save**. In our example, we leave the fallback as Deny.
13. Click the + symbol between **LDAP Auth** and **Allow**.
14. In the General Purpose section, click the **Resource Assign** option button, and then click **Add Item**.
15. Click the **Add new entry** button.
16. Click **Set Network Access Source**, and then click the option button for the Network Access Source you created in *To configure remote access*, on page 6. In our example, we click **London_Remote_Access**. This associates the Lease Pool and other settings. Click the **Update** button. You return to the Resource Assign page.
17. Click **Set Webtop**, and then click the option button for the Webtop you created in *Creating a Webtop*, on page 10. In our example, we click **London_Webtop**. Click the **Update** button.
18. Click the **Save** button. The Resource Assignment window closes and you return to the Visual Policy Editor main page.
At this point you have the basics for a functional access policy.
19. Click the yellow **Apply Access Policy** link in the upper left part of the window. You always have to apply an access policy before it takes effect.

20. Click the **Close** button on the upper right to close the VPE.

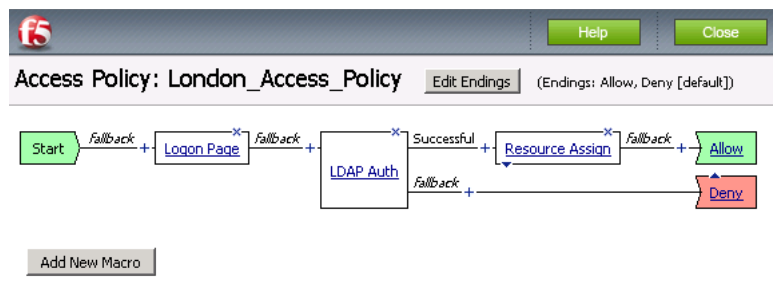


Figure 8 Final result of our Access Policy example

Configuring network entry points

The next task is to create the external virtual server that allows users to initiate their connection to the SSL VPN from either the web browser or the BIG-IP Edge Client for Windows. In our example, we have chosen to allow DTLS as a connection method and we will create two virtual servers, one for TCP 443 and one for UDP 4433.

The first task is to create profiles that are used by the virtual servers.

Creating the profiles

The next step is to create the profiles. Although you may use the default profiles, we strongly recommend you create new profiles based on the default parent profiles. By creating new profiles, you may easily modify the profile settings specific to your deployment without altering default global behaviors.

Creating TCP profiles

The next task is to create the TCP profiles. We recommend creating **tcp-lan-optimized** and **tcp-wan-optimized** profiles.

Creating the LAN optimized TCP profile

The first TCP profile we create is the LAN optimized profile.

To create a new LAN optimized TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens by default.
2. On the Menu bar, from the **Protocol** menu, select **TCP**.
3. Click the **Create** button. The New TCP Profile screen opens.

-
4. In the **Name** box, type a name for this profile. In our example, we type **edge_tcp_lan**.
 5. Modify any of the settings as applicable for your network. See the online help for more information on the configuration options. In our example, we leave the settings at their default levels.
 6. Click the **Finished** button.

Creating the WAN optimized TCP profile

The next task is to create the WAN optimized profile.

To create a new WAN optimized TCP profile

1. On the Main tab, expand **Local Traffic**, click **Profiles**, and then, on the Menu bar, from the **Protocol** menu, select **TCP**.
2. Click the **Create** button. The New TCP Profile screen opens.
3. In the **Name** box, type a name for this profile. In our example, we type **edge_tcp_wan**.
4. Modify any of the settings as applicable for your network. See the online help for more information on the configuration options. In our example, we leave the settings at their default levels.
5. Click the **Finished** button.

Creating the HTTP profile

The next profile to create is the HTTP profile. This profile is required for the VPN to function. This should be a simple HTTP profile with no optimization (compression or caching).

To create the HTTP profile

1. On the Main tab, expand **Local Traffic**, click **Profiles**, and then click the **Create** button.
2. In the **Name** box, type a name for this profile. In our example, we type **edge-http**.
3. Modify any of the settings as applicable for your network, but **do not** enable compression or RAM Cache (both are disabled by default when using the **http** parent profile). See the online help for more information on the configuration options. In our example, we leave the settings at their default levels.
4. Click the **Finished** button.

Creating a Client SSL profile

The next step is to create an SSL profile. This profile contains SSL certificate and Key information for offloading SSL traffic. The first task is to import the certificate and key (for this Deployment Guide, we assume that you already have obtained the required SSL certificates, but they are not yet installed on the BIG-IP LTM system. If you do not have a certificate and key, see the BIG-IP documentation).

To import a key or certificate

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

The next task is to create the SSL profile that uses the certificate and key you just imported.

To create a new Client SSL profile

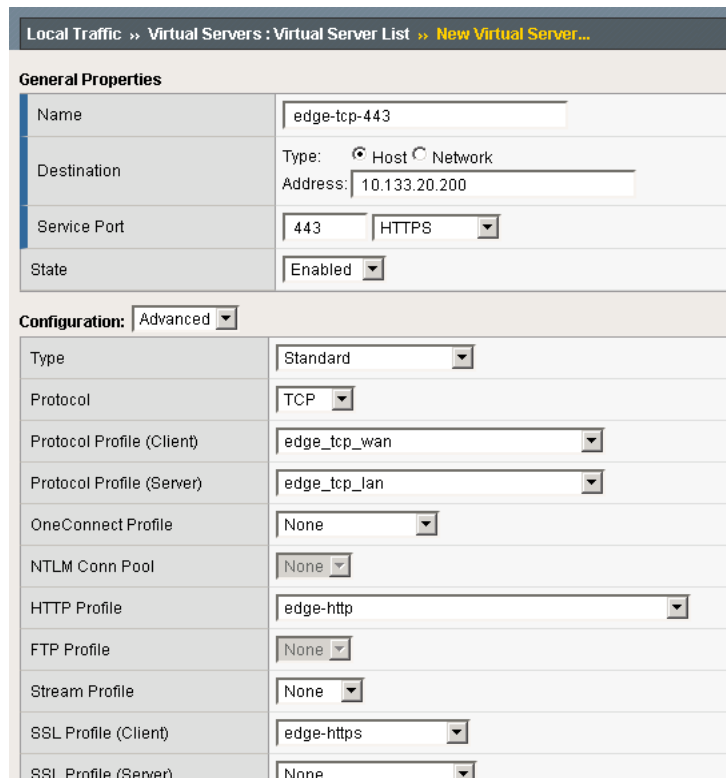
1. On the Main tab, expand **Local Traffic**, click **Profiles**, and then, on the Menu bar, from the **SSL** menu, select **Client**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this profile. In our example, we type **edge_https**.
4. In the Configuration section, click a check in the **Certificate** and **Key** Custom boxes.
5. From the **Certificate** list, select the name of the Certificate you imported in the *Importing keys and certificates* section.
6. From the **Key** list, select the key you imported in the *Importing keys and certificates* section.
7. Click the **Finished** button.

Creating the virtual servers

The next task is to create the virtual servers for TCP 443 and UDP 4433.

To create the virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. Click the **Create** button. The New Virtual Server screen opens.
3. In the **Name** box, type a name. We type **edge-tcp-443**.
4. In the **Destination** section, select the **Host** option button.
5. In the **Address** box, type the IP address. We use **10.133.20.200**.
6. In the **Service Port** box, type **443**, or select **HTTPS** from the list.
7. In the Configuration section, select **Advanced** from the list.
8. From the **Protocol Profile (Client)** list, select the profile you created in *Creating the WAN optimized TCP profile*, on page 15. In our example, we select **edge_tcp_wan**. This is optional.
9. From the **Protocol Profile (Server)** list, select the profile you created in *Creating the LAN optimized TCP profile*, on page 14. In our example, we select **edge_tcp_lan**.
10. From the **HTTP Profile** list, select the profile you created in *Creating the HTTP profile*. In our example, we select **edge-http**.
11. From the **SSL Profile (Client)** list, select the SSL profile you created in *Creating a Client SSL profile*. We select **edge_https**.



Local Traffic » Virtual Servers : Virtual Server List » New Virtual Server...	
General Properties	
Name	edge-tcp-443
Destination	Type: <input checked="" type="radio"/> Host <input type="radio"/> Network Address: 10.133.20.200
Service Port	443 HTTPS
State	Enabled
Configuration: Advanced	
Type	Standard
Protocol	TCP
Protocol Profile (Client)	edge_tcp_wan
Protocol Profile (Server)	edge_tcp_lan
OneConnect Profile	None
NTLM Conn Pool	None
HTTP Profile	edge-http
FTP Profile	None
Stream Profile	None
SSL Profile (Client)	edge-https
SSL Profile (Server)	None

Figure 9 New virtual server configuration (truncated)

12. In the Access Policy section, from the **Access Profile** list, select the name of the policy you created in *Creating an Access Profile*, on page 11. In our example, we select **London_Access_Policy**.
13. From the **Connectivity Profile** list, select the profile you created in *Creating a Connectivity Profile*, on page 10. In our example, we select **London_Connectivity_Profile**.
14. Leave the **Rewrite Profile** list set to **None**.

Access Policy	
Access Profile	London_Access_Policy ▾
Connectivity Profile	London_Connectivity_Profile ▾
Rewrite Profile	None ▾

Figure 10 Access Policy section of the virtual server configuration

15. **Do not** configure any of the options in the WAN Optimization section.
16. Click the **Finished** button (this virtual server does not have any Resources).
17. Repeat this entire procedure for the UDP virtual server with the following exceptions.
 - In Step 3, give this virtual server a unique name.
 - In Step 5, use the appropriate IP address.
 - In Step 6, in the **Service Port** box, type **4433**.
 - After Step 7, from the **Protocol** list, select **UDP**.
 - All other settings are the same.

Configuring site-to-site WAN optimization for application access across data centers

In this section, we configure the site-to-site WAN optimization for application access across data centers scenario with the WAN Optimization Module component of BIG-IP Edge Gateway. The WAN Optimization Module accelerates application traffic between datacenters. In our example the remote access function is provided centrally and applications are distributed among datacenters. Much of the configuration in this section is completed on both the local and remote BIG-IP systems.

Prerequisites and configuration notes

The following are notes specific to the WAN Optimization module.

- ◆ This configuration in this section is written with the assumption that application servers in each location are presented in the DMZ at each location. BIG-IP WOM will not be used to cross security layer boundaries.
- ◆ The BIG-IP Edge Gateway should not advertise its local internal application server network. Rather the BIG-IP Edge Gateway at each location should advertise the DMZ subnet addresses.
- ◆ Each DMZ should route to its local Edge Gateway to access remote DMZ and client access subnets, if not using SNAT.

Configuring the WAN Optimization module

Use the following procedures to configure the WAN Optimization module.

Creating the iSession profile

In this procedure, we create an iSession profile. The iSession profile tells the system how you want to optimize traffic; encryption, deduplication and compression are controlled in the iSession profile. Creating a custom iSession profile for each application is a best practice.

To create the iSession profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, from the **Services** menu, select **iSession**.
3. Click the **Create** button.
4. In the **Name** box, type a name for this profile. In our example, we type **edge-isession**.

5. Click the **Custom** box for **Reuse Connection**, and then select **Disabled** from the list.
6. Leave the other settings at the default levels.
7. Click the **Finished** button.

Local Traffic >> Profiles : Services : iSession >> New iSession Profile...

General Properties

Name	edge-isession
Parent Profile	isession

Settings Custom

Mode	Enabled	<input type="checkbox"/>
Deduplication	Enabled	<input type="checkbox"/>
Reuse Connection	Disabled	<input checked="" type="checkbox"/>
Target Virtual	match all	<input type="checkbox"/>
Port Transparency	Enabled	<input type="checkbox"/>

Compression Settings Custom

Adaptive	Enabled	<input type="checkbox"/>
Deflate	Enabled	<input type="checkbox"/>
Deflate Level	1	<input type="checkbox"/>
LZO	Enabled	<input type="checkbox"/>
Null	Enabled	<input type="checkbox"/>

Cancel Repeat Finished

Figure 11 iSession profile configuration

Configuring BIG-IP WOM using the Quick Start template

The next procedure uses the Quick Start Template to configure the WAN Optimization module.

Note: The Quick Start Template is only available in BIG-IP WOM version 10.2 or later.

To configure the WAN Optimization Module using the Quick Start template

1. On the Main tab of the local BIG-IP system, expand **WAN Optimization**, and then click **Quick Start**. The Quick Start configuration screen opens.
2. In the **WAN Self IP Address** box, type the BIG-IP self IP address you provisioned for iSession endpoint in this data center.
3. From the **Discovery** list, select **Enabled**.

-
4. In the **Select VLANs** section, select the VLAN you configured and place it in the LAN and WAN VLAN boxes.
 5. Choose the iSession profile you created above
 6. For authentication and Encryption we recommend that you use the default profiles.
 7. From the **Create Optimized Applications** list, select the applications and protocols you wish to accelerate between datacenters. If your application protocol is not on this list you may create a new application optimization policy by clicking **Click this link to create other optimized applications**.
 8. Click **Apply**.

The next part of this procedure is to configure the Advertised routes

9. Configure the Advertised Routes in datacenter 1:
 - a) On the Main tab, expand **WAN Optimization**, and then click **Advertised Routes**.
 - b) Click the **Create** button.
 - c) In the **Address** box, type the local DMZ subnet address for this datacenter.
 - d) In the **Netmask** box, type the corresponding subnet mask.
 - e) You can optionally type a descriptive Label.
 - f) Leave the Mode set to **Included**.
 - g) If you are not using SNAT for remote access clients, create another advertised route for the network access subnet you created above.
10. Configure Advertised Routes in datacenter 2:
 - a) On the Main tab, expand **WAN Optimization**, and then click **Advertised Routes**.
 - b) Click the **Create** button.
 - c) In the **Address** box, type the local DMZ subnet address for this datacenter.
 - d) In the **Netmask** box, type the corresponding subnet mask.
 - e) You can optionally type a descriptive Label.
 - f) Leave the Mode set to **Included**.
 - g) Click **Finished**.

 **Important**

Repeat this procedure in each data center.

After performing this procedure on both BIG-IP systems, continue with the following procedure, in which you connect your two BIG-IP systems together via an iSession tunnel by identifying each remote endpoint. If dynamic discovery was left on (as in step 3), you will only perform the following procedure on one of the BIG-IP systems. If you did not, you must repeat this procedure on the remote BIG-IP system.

To configure the remote endpoints in data center 1

1. On the Menu bar, click **Remote Endpoints**. The remote endpoints configuration screen appears.
2. Click the **Create** button.
3. In the **IP Address** box, type the IP Address of the remote endpoint iSession Self IP Address for the BIG-IP Edge Gateway in Datacenter 2.
4. Accept default values.
5. Click **Finished**.

Configure Remote Endpoints in datacenter 2

1. On the menu bar, click **Remote Endpoints**.
2. Click the **Create** button.
3. In the **IP Address** box, type the IP Address of the remote endpoint iSession Self IP Address for the BIG-IP Edge Gateway in Datacenter 1.
4. Accept default values,
5. Click **Finished**.

Configuring authenticated access to accelerated web applications

In this section, we configure the authenticated access to accelerated web applications scenario, using the BIG-IP APM to authenticate users and then connect them to an accelerated web application. This ensures that application connections coming into your DMZ are authenticated before being passed onto the application. You can optionally require prelogin checks.

For the configuration in this section, you must have a license for BIG-IP LTM, APM and WebAccelerator. For more information on licensing options, talk to your sales representative.

Creating an AAA server

The first task is to create an AAA server. To create the AAA server, use the procedure *Creating an AAA resource*, on page 8. Give the AAA server a unique name. Configure all other settings as applicable for your implementation.

Creating the SSO configuration

The next task is to create a Single Sign-On Configuration that defines the credentials that are cached.

To create the SSO configuration

1. On the Main tab, expand **Access Policy**, and then click **SSO Configurations**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this profile. In our example, we type **SanJose-SSO**.
4. From the **SSO Method** list, select the appropriate SSO method. In our example, we select **HTTP Basic**.
5. In the **Username Source** box, type the user name source. In our example, we leave the default: **session.sso.token.last.username**.
6. In the **Password Source** box, type the user name source. In our example, we leave the default: **session.sso.token.last.password**.
7. From the **Access Management Method** list, select **None**.
8. Click **Finished**.

Access Policy >> SSO Configurations >> New SSO Configuration...	
General Properties	
Name	SanJose-SSO
SSO Method	None
SSO Method Configuration	
Username Source	session.sso.token.last.username
Password Source	session.sso.token.last.password
External Access Management	
Access Management Method	None
Cancel Finished	

Figure 12 SSO configuration

Creating an Access Profile

Next, we create an Access profile.

To create an Access Profile

1. On the Main tab, expand **Access Policy**, and then click **Access Profiles**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this profile. In our example, we type **SanJose_Access_Policy**.
4. In the Settings section, configure the options as applicable for your configuration. In our example, we leave all of the settings at their defaults. Note that depending on licensing, the number of concurrent users may be limited. The other timeouts are administrative choices.
5. In the Configuration section, from the SSO Configuration list, select the SSO configuration you created in *Creating the SSO configuration*, on page 23.
6. Configure the rest of the settings in the Configuration as applicable to your environment. In our example, we accept all of the defaults. However, we do leave **Secure Cookie** checked.
7. In the Language Settings section, if you are configuring the Edge Gateway in a language other than English, configure as applicable for your language. In our example, we accept English as the default language.
8. Click **Finished**.

Editing the Access Profile with the Visual Policy Editor

The next task is to edit the Access Policy using the Visual Policy Editor (VPE). The VPE is a powerful visual scripting language that offers virtually unlimited options in configuring an Access Policy. For detailed information on the VPE please see the product documentation.

The following is just an example, you can add more or different prelogon checks as applicable to your configuration.

To edit the Access Profile

1. On the Main tab, expand **Access Policy**, and then click **Access Profiles**.
2. Locate the Access Profile you just created, and in the Access Policy column, click **Edit**.
The Visual Policy Editor opens in a new window.
3. Click the + symbol between **Start** and **Deny**. A box opens with options for different actions.
4. Click the **Logon Page** option button, and then click the **Add Item** button at the bottom of the box.
5. Configure the Properties as applicable for your configuration. In our example, we leave the settings at the defaults.
6. Click the **Save** button.
7. Click the + symbol between **Logon Page** and **Deny**.
8. In the Authentication section, click the **LDAP Auth** option button, and then click the **Add Item** button.
9. From the Server list, select the AAA Source you created in *Creating an AAA server*, on page 23.
10. Add **SearchDN** and **SearchFilter** items as applicable for your configuration.
11. Click the **Save** button. You now see two paths, **Successful** and **Fall Back**.
12. Click the **Deny** box from the path leading from Successful. The Select Ending box opens.
13. Click the **Allow** button, and then click **Save**. In our example, we leave the fallback as Deny.
14. Click the yellow **Apply Access Policy** link in the upper left part of the window. You always have to apply an access policy before it takes effect.
15. Click the **Close** button on the upper right to close the VPE.

Creating the profiles

The next step is to create the profiles. Although you may use the default profiles, we strongly recommend you create new profiles based on the default parent profiles. By creating new profiles, you may easily modify the profile settings specific to your deployment without altering default global behaviors.

Creating TCP profiles

The next task is to create the TCP profiles. We recommend creating **tcp-lan-optimized** and **tcp-wan-optimized** profiles.

Creating the LAN optimized TCP profile

The first TCP profile we create is the LAN optimized profile.

To create a new LAN optimized TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens by default.
2. On the Menu bar, from the **Protocol** menu, select **TCP**.
3. Click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **edge_tcp_lan**.
5. Modify any of the settings as applicable for your network. See the online help for more information on the configuration options. In our example, we leave the settings at their default levels.
6. Click the **Finished** button.

Creating the WAN optimized TCP profile

The next task is to create the WAN optimized profile.

To create a new WAN optimized TCP profile

1. On the Main tab, expand **Local Traffic**, click **Profiles**, and then, on the Menu bar, from the **Protocol** menu, select **TCP**.
2. Click the **Create** button. The New TCP Profile screen opens.
3. In the **Name** box, type a name for this profile. In our example, we type **edge_tcp_wan**.
4. Modify any of the settings as applicable for your network. See the online help for more information on the configuration options. In our example, we leave the settings at their default levels.
5. Click the **Finished** button.

Creating the HTTP profile

The next profile to create is the HTTP profile. This profile is required for the VPN to function. This should be a simple HTTP profile with no optimization (compression or caching).

To create the HTTP profile

1. On the Main tab, expand **Local Traffic**, click **Profiles**, and then click the **Create** button.
2. In the **Name** box, type a name for this profile. In our example, we type **edge-http**.
3. Modify any of the settings as applicable for your network, but **do not** enable compression or RAM Cache. See the online help for more information on the configuration options. In our example, we leave the settings at their default levels.
4. Click the **Finished** button.

Creating a Client SSL profile

If your configuration requires offloading SSL transactions on the BIG-IP, the next task is to create an SSL profile. This profile contains SSL certificate and Key information for offloading SSL traffic.

The first task is to import the certificate and key, see *To import a key or certificate*, on page 16.

To create a new Client SSL profile

1. On the Main tab, expand **Local Traffic**, click **Profiles**, and then, on the Menu bar, from the **SSL** menu, select **Client**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this profile. In our example, we type **edge_https**.
4. In the Configuration section, click a check in the **Certificate** and **Key Custom** boxes.
5. From the **Certificate** list, select the name of the Certificate you imported in the *Importing keys and certificates* section.
6. From the **Key** list, select the key you imported in the *Importing keys and certificates* section.
7. Click the **Finished** button.

Creating the virtual server

The next task is to create the virtual server.

To create the virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. Click the **Create** button. The New Virtual Server screen opens.
3. In the **Name** box, type a name for this virtual server. In our example, we type **SanJose-vs**.
4. In the **Destination** section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we use **10.133.20.200**.
6. In the **Service Port** box, type **443**, or select **HTTPS** from the list.
7. In the Configuration section, select **Advanced** from the list. The Advanced configuration options appear.
8. From the **Protocol Profile (Client)** list, select the name of the profile you created in *Creating the WAN optimized TCP profile*, on page 26. In our example, we select **edge_tcp_wan**. This is optional.
9. From the **Protocol Profile (Server)** list, select the name of the profile you created in *Creating the LAN optimized TCP profile*, on page 26. In our example, we select **edge_tcp_lan**.
10. From the **HTTP Profile** list, select the name of the profile you created in the *Creating the HTTP profile* section. In our example, we select **edge-http**.
11. If you created an optional SSL profile, from the **SSL Profile (Client)** list, select the SSL profile you created in the *Creating a Client SSL profile* section. In our example, we select **edge_https**.
12. In the Access Policy section, from the **Access Profile** list, select the name of the policy you created in *Creating an Access Profile*, on page 24. In our example, we select **SanJose_Access_Policy**.
13. Leave the **Rewrite Profile** list set to **None**.
14. **Do not** configure any of the options in the WAN Optimization section.
15. Click the **Finished** button (this virtual server does not have any Resources).

Configuring the WebAccelerator

In this section, we configure the WebAccelerator to accelerate application traffic.

For this configuration, we assume you have already created the necessary BIG-IP configuration objects (monitor, pool, profiles, virtual server) for your application. If you have not created these objects, see the F5 deployment guide appropriate for your configuration (<http://www.f5.com/solutions/resources/deployment-guides.html>).

Creating an HTTP class

In this procedure, we create an HTTP class with Web Acceleration turned on.

To create the HTTP class

1. On the Main tab, expand **WebAccelerator**, and then click **Class Profiles**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this class. In our example, we type **application-class**.
4. From the **WebAccelerator** list, make sure **Enabled** is selected.
5. In the Hosts row, from the list select **Match Only**. The Host List options appear.
 - a) In the **Host** box, type the host name that your end users use to access your application. In our example, we type **sharepoint.example.com**.
 - b) Leave the Entry Type at **Pattern String**.
 - c) Click the **Add** button.
 - d) Repeat these sub-steps for any other host names users might use to access the application.
6. The rest of the settings are optional, configure them as applicable for your deployment.
7. Click **Finished**.

Configuring a WebAccelerator application

The next procedure is to create a WebAccelerator Application for your particular application. The Application provides key information to the WebAccelerator so that it can handle requests to your application appropriately.

To create a new Application

1. On the Main tab, expand **WebAccelerator**, and then click **Applications**. The Application screen of the WebAccelerator UI opens in a new window.
2. Click the **New Application** button.
3. In the Application Name box, type a name for your application. In our example, we type **sharepoint-application**.
4. In the **Description** box, you can optionally type a description.
5. From the **Local Policies** list, select the appropriate application from the list. In our example, we select **Microsoft SharePoint Services 2007**.
6. In the **Requested Host** box, type the host name that your end users use to access the application. This should be the same host name you used in Step 5a in the preceding procedure. In our example, we type **sharepoint.example.com**.

If you have additional host names, click the **Add Host** button and enter the host name(s).

7. Click the **Save** button.

Modifying the virtual server

The next task is to modify the existing virtual server for your application to use the HTTP class you created.

To modify the Virtual Server to use the Class profile

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. From the **Virtual Server** list, click the name of the virtual server you created for your application. The General Properties screen for the Virtual Server opens.
3. On the Menu bar, click **Resources**.
4. In the HTTP Class Profiles section, click the **Manage** button.
5. From the **Available** list, select the name of the HTTP Class Profile you created in *Creating an HTTP class*, on page 29, and click the Add (<<) button to move it to the Enabled box. In our example, we select **application-class**.
6. Click the **Finished** button. The HTTP Class Profile is now associated with the Virtual Server.

If you are creating a new virtual server for your application, be sure to associate the HTTP Class profile with virtual server.