



Deploying the BIG-IP Edge Gateway and Local Traffic Manager with VMware View 4 and 4.5

Table of Contents

Deploying F5 with VMware View 4.5

Product versions and revision history	1-1
Configuration flow	1-2

Configuring the BIG-IP Edge Gateway

Configuring remote access	1-4
Creating a Connectivity Profile	1-6
Creating a Webtop	1-7
Creating an AAA Server	1-7
Creating a Web Application	1-8
Creating an Access Profile	1-9
Editing the Access Profile with the Visual Policy Editor	1-10
Creating the Network Access virtual server configuration objects	1-15
Creating the profiles	1-15
Creating the iRule	1-17
Creating the virtual servers	1-18

Configuring the BIG-IP LTM for VMware View 4.5

Prerequisites and configuration notes	2-1
Modifying the VMware Virtual Desktop Manager global settings	2-2
Configuring the External URL	2-2
Configuring the BIG-IP LTM system for VMware Connection Brokers	2-3
Creating the health monitor	2-3
Creating the View Manager server pool	2-3
Creating the Universal Inspection Engine persistence iRule	2-5
Using SSL certificates and keys	2-6
Creating BIG-IP LTM profiles	2-7
Creating the virtual server	2-11



I

Deploying the BIG-IP Edge Gateway for VMware View

Deploying F5 with VMware View 4.5

Welcome to the F5 Deployment Guide for VMware View (formerly Virtual Desktop Infrastructure: VDI). This document provides guidance and configuration procedures for deploying the BIG-IP Local Traffic Manager (LTM) and BIG-IP Edge Gateway with VMware View.

The VMware View portfolio of products lets IT run virtual desktops in the datacenter while giving end users a single view of all their applications and data in a familiar, personalized environment on any device at any location.

One of the unique features of this deployment is the ability of the BIG-IP LTM system to persist client to broker connections on a session by session basis. Other implementations commonly use simple/source address persistence, where all the connections from a single IP address are sent to one server. With the iRule described later in this document, the BIG-IP LTM is able to direct traffic with greater precision, resulting in a more uniform load distribution on the connection servers.

The BIG-IP Edge Gateway provides pre-logon checks to the endpoint device and supports a broad range of authentication mechanisms, including two-factor schemes and various back-end directory services. Edge Gateway can also enforce Active Directory group policies on corporate-owned and non-corporate-owned assets during the duration of the connection. Additionally, once authenticated, Edge Gateway guarantees the encryption of all VMware View transport protocols, whether natively encrypted or not. With all these features, Edge Gateway is able to replace the View Security Server.

This guide is broken into two main sections:

- *Configuring the BIG-IP Edge Gateway*, on page 1-4
- *Configuring the BIG-IP LTM for VMware View 4.5*, on page 2-1

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

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

Product versions and revision history

Product and versions tested for this deployment guide:

Product Tested	Version Tested
BIG-IP Edge Gateway	10.2
BIG-IP LTM	10.2
VMware View	4.5

Document Version	Description
1.0	New guide for View 4.0
2.0	Added BIG-IP Edge Gateway chapter
2.1	Modified line 15 of the iRule on page 1-17 from set password \$value to set password [URI::decode \$value] to support the full range of characters included in RFC2396.
3.0	Added support for View 4.5

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.

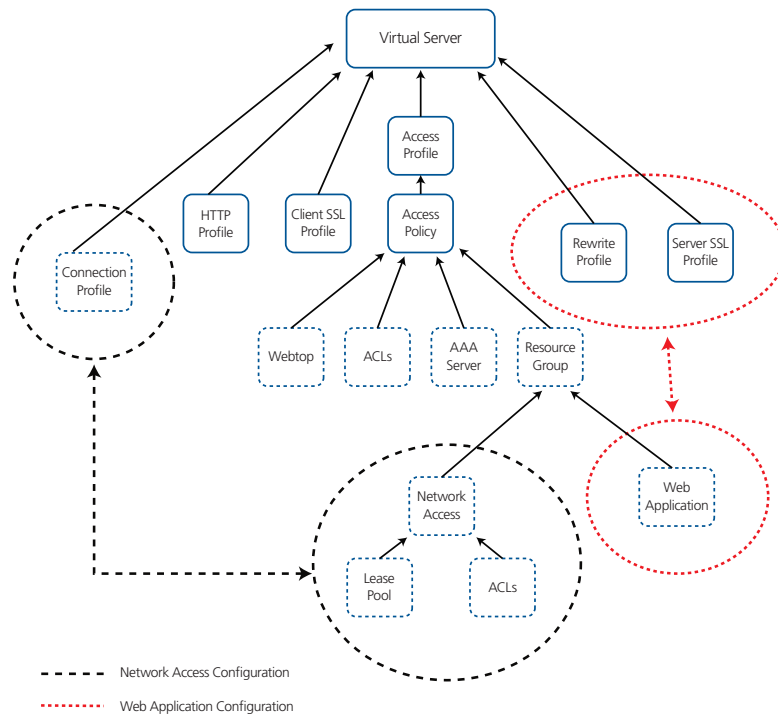


Figure 1.1 Configuration flow

Figure is a logical configuration example of this deployment.

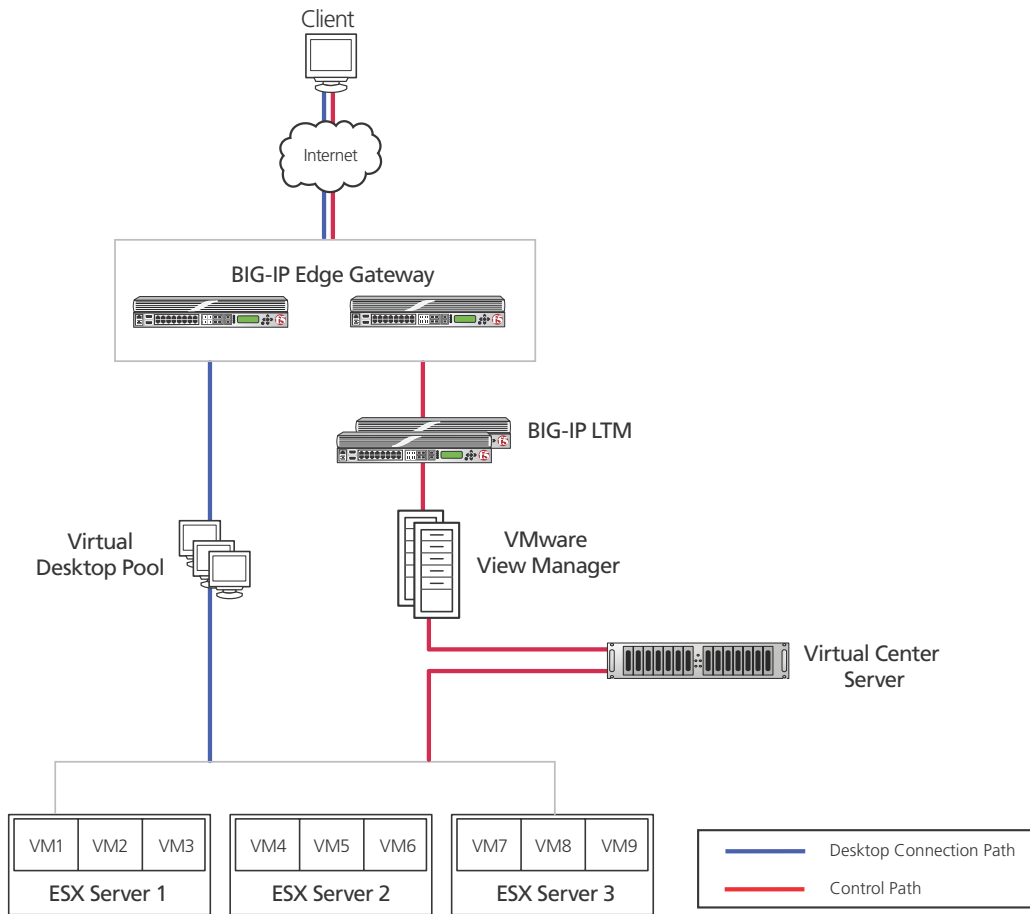


Figure 1.2 Logical configuration example

Configuring the BIG-IP Edge Gateway

Use the following procedures to configure the BIG-IP Edge Gateway for VMware View 4.5.

Configuring remote access

To configure Remote Access, 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.

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 **View-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 **View-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 **192.0.2.1** to **192.0.2.255**.
 - d) Click the **Add** button.

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

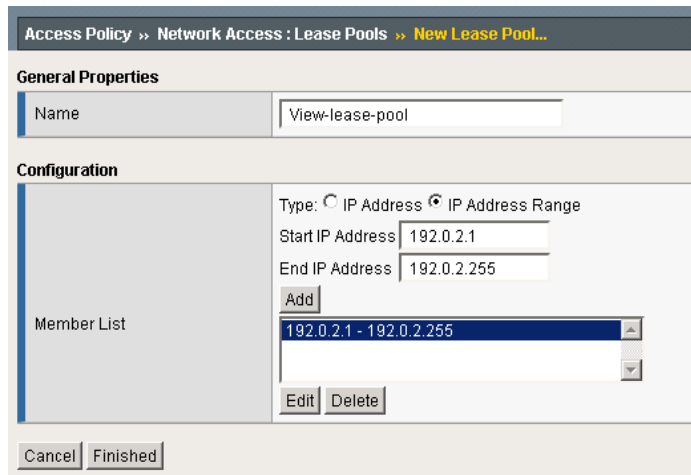


Figure 1.3 Lease Pool configuration

5. If necessary, from the **Lease Pool** list, select the lease pool you just created. In our example, we select **View-lease-pool**.
6. From the **Compression** list, select **GZIP Compression**. This allows both the web browser client and the thick client to take advantage of compression between the client and the remote access server.

Note: If Datagram TLS (DTLS) is configured (UDP based communication between client and Remote Access Server) GZIP compression is automatically disabled. DTLS and GZIP compression are incompatible with one another. If you enable GZIP compression it will be used for TCP connections. DTLS clients will use compression for network access tunnels.

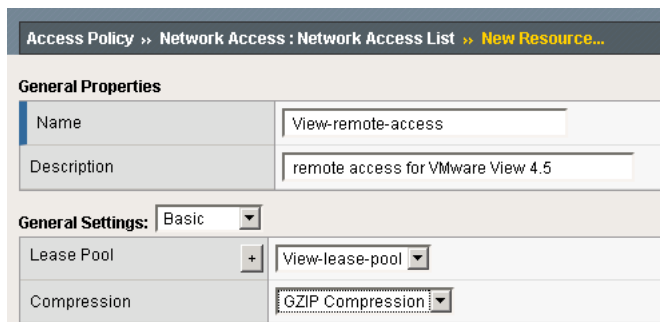


Figure 1.4 Network Access General settings

7. From the **Client Settings** list, select **Advanced**.

8. 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 9.
 - 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, in the **IP address** and **Mask** boxes, type the IP address and Mask of the LAN Address space that should go through the tunnel. In our example we indicate that the LAN address space is **192.168.0.0/16**.
 - In the DNS Address Space section, in the **DNS** box, type the DNS suffixes 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 the LAN should be inaccessible to remote access clients, it can be entered here. In our example, we indicate that **192.168.10.0/24** is excluded.
9. 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 18). If clients cannot connect with DTLS, they fall back to TCP based SSL.

10. Click **Finished**.

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 **View-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 BIG-IP Edge, a Network Webtop is a pointer that initiates the download of the Edge client for browsers.

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 **View-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	View-webtop
Type	Network Access
Configuration	
Minimize To Tray	<input checked="" type="checkbox"/> Enabled
Cancel Repeat Finished	

Figure 1.5 New Webtop

Creating an AAA Server

The Edge Gateway does not have a built-in authentication store therefore an authentication source must be specified. In this procedure, we create an AAA server.

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 **View-ActiveDirectory**.
4. From the **Type** list, select the appropriate authentication method. For this example, we select **Active Directory**.
5. In the Configuration section, type the appropriate information relevant to your Active Directory services.
6. Click **Finished**.

Creating a Web Application

The next task is to create a Web Application. This Web Application contains the IP address of the BIG-IP LTM virtual server for the Connection Broker servers, where users are directed if the prelogon policy cannot detect the View client.

To create a Web Application

1. On the Main tab, expand **Access Policy**, and then click **Web Applications**.
2. Click the **Create** button.
3. In the **Name** box, type a name. In our example, we type **DownloadViewClient**.
4. In the **Patching** section, from the **Type** list, select **Minimal Patching**, and then click the **Scheme Patching** box.
5. Click the **Create** button. The Resource Items appear.
6. Click the **Add** button to the right of Resource Items.
7. In the Destination row, click the **IP Address** option button, and then in the **IP Address** box, type the IP address of the BIG-IP LTM virtual server you created for the Connection Broker servers in *Creating the virtual server*, on page 2-11.
8. In the **Port** box, type **80**.
9. From the **Scheme** list, select **HTTP**.
10. In the **Paths** box, type **/***.
11. From the **Compression** list, select **GZIP Compression**.
12. Leave the other settings at their defaults.
13. Click the **Finished** button.

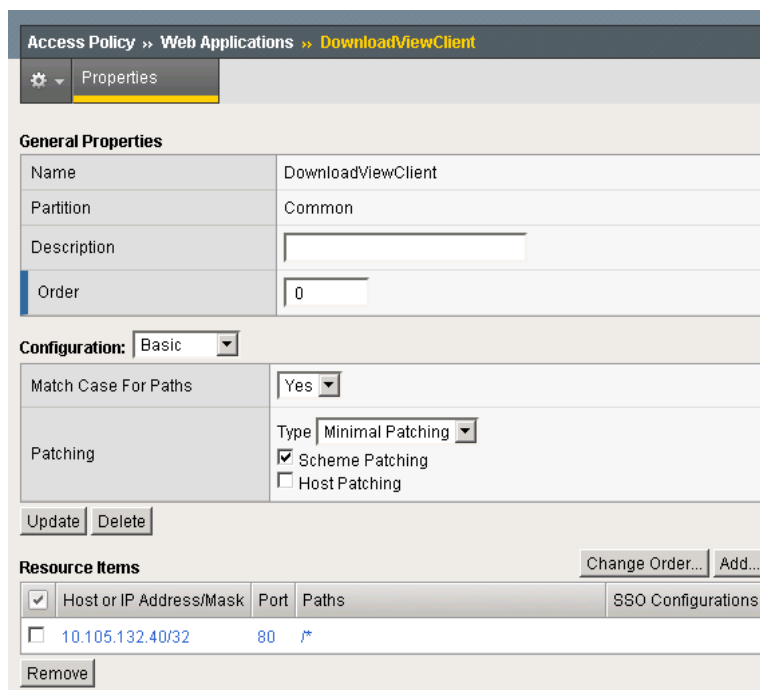


Figure 1.6 Web Application configuration

Creating an Access Profile

The Access Profile ties together all of the other pieces in order to create a Network Connection VPN Tunnel. The Access Profile is also where the Visual Policy Editor (VPE) is located, which allows for complex workflows to be designed.

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 **View-access**.
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 View-access profile and edit it 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. In this example, we first check the client's operating system, send an Unsupported Operating System to any user who isn't using Microsoft Windows XP.

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. In the Server Side Checks section at the bottom of the box, click the Client OS option button, and then click the **Add Item** button at the bottom of the box. Paths for eight different operating systems appear.
5. Click the **Add New Macro** button. The new macro box opens.
 - a) In the **Name** box, type a name for this macro. In our example, we type **UnsupportedOSMessage**.
 - b) Click the Save button. The Macro appears under the Access Policy.
 - c) Click the Expand (+) button next to UnsupportedOSMessage.
 - d) Click the + symbol between **In** and **Out**. A box opens with options for different actions.
 - e) Click the **Message box** option button, and then click **Add Item**.
 - f) In the **Name** box, type a unique name for this box. In our example, we type **serviceNotAvailableforThisOS**.
 - g) You can optionally change the Language.

-
- h) In the **Message** box, type the message you want users to see. In our example, we type **This service is available for Windows XP clients only**.
 - i) You can optionally modify the Link text. Clicking the link sends the user to the next object in the path, which is Deny in our example.
 - j) Click the **Save** button. The macro is now ready to use in the following step.
 6. Click the **+** symbol between **Windows 7** and **Deny**. A box opens with options for different actions.
 7. In the Macrocalls section, click the option button for the macro you just created, and then click the **Add Item** button. In our example, we click **UnsupportedOSMessage**.
 8. Repeat steps 6 and 7 for each of the operating systems you want to deny. In our example, we repeat these steps for all operating systems except Windows XP.
 9. Click the **+** symbol between **Windows XP** and **Deny**.
 10. In the General Purpose section, click the **Logon Page** option button, and then click **Add Item**.
 11. Modify any settings as applicable for your configuration. In our example, we leave these at the defaults. Click **Save**.
 12. Click the **+** symbol between **Logon Page** and **Deny**.
 13. In the Authentication section, click the **AD Auth** option button, and then click the **Add Item** button. The AD Auth box opens. Complete the following:
 - a) In the **Name** box, you can optionally type a new name. In our example, we type **AuthenticatedUser**.
 - b) From the **Server** list, select the name of the AAA server you created in *Creating an AAA Server*, on page 1-7. In our example, we select **View-ActiveDirectory**.
 - c) Modify the remaining settings as applicable for your configuration. In our example, we leave the defaults.
 - d) Click **Save**.
 14. On the Successful path between **AuthenticatedUser** and **Deny**, click the **+** symbol.
 15. In the General Purpose section, click the **iRule Event** option button, and then click the **Add Item** button. The iRule Event page opens. Complete the following:
 - a) In the **Name** box, you can optionally type a new name. In our example, we type **CopyPasswordToSessionVar**.

- b) In the **ID** box, type a numeric ID for this event. In our example, we type **1**. If you are using this agent elsewhere, make sure this ID is unique.
 - c) Click **Save**.
16. Click the **+** symbol between **CopyPasswordToSessionVar** and **Deny**.
 17. In the Client Side Check section, click the **Windows File Check** option button, and then click the **Add Item** button. The Windows File Checker page opens. Complete the following:
 - a) In the **Name** box, you can optionally type a new name. In our example, we type **checkForViewClient**.
 - a) Click the **Add new entry** button.
 - b) In the **FileName** box, type the path to the View client as appropriate for your View deployment. In our example, we type:
`C:\\Program Files\\VMware\\VMware View\\Client\\bin\\wswc.exe`

Note: The double backslashes are required for the inspector to check for the file. If your View client is installed in a custom location be sure to set the correct path to the executable.

 - c) Leave the rest of the settings at their default levels.
 - d) Click the **Save** button.
 18. On the Successful path between **checkForViewClient** and **Deny**, click the **+** symbol.
 19. In the General Purpose section, click the **Resource Assign** option button, and then click the **Add Item** button. The Resource Assignment page opens. Complete the following:
 - a) In the **Name** box, you can optionally type a name. In our example, we type **AssignViewClient**.
 - b) Click the **Add new entry** button.
 - c) Click the **Set Network Access Resource** link, and then click the option button for the Network Access resource you created in *Configuring remote access*, on page 1-4. In our example, we click **View-remote-access**. Click **Update**.
 - d) Click the **Set Webtop** link, and then click the option button for the Webtop you created in *Creating a Webtop*, on page 1-7. In our example, we click **View-webtop**. Click **Update**.
 - e) Click the **Save** button.
 20. On the Fallback path between **AssignViewClient** and **Deny**, click the **+** symbol.

-
21. In the General Purpose section, click the **Variable Assign** option button, and then click the **Add Item** button. The Resource Assignment page opens. Complete the following:
 - a) In the **Name** box, you can optionally type a new name. In our example, we type **configureViewSSO**.
 - b) Click the **Add new entry** button.
 - c) Click the **change** button.
 - d) From the list on the left, select **Configuration Variable**.
 - e) From the **Property** list, select **application launch**.
 - f) In the Custom Expression box on the right, copy and paste the following expression, replacing the relevant information (see note following):

```
expr {
"<application_launch><item><path>C:\\Program Files\\VMware\\VMware
View\\Client\\bin\\wswc.exe</path><parameter>-username [mcget
{session.logon.last.username}] -password [mcget {session.vmware.sso}] -domainName BD
-serverURL
http://10.105.132.40:80</parameter><os_type>WINDOWS</os_type></item></application_launch>" }
```

Note: If your View client is installed in a custom location be sure to set the correct path to the executable. Our Domain is "BD", be sure to insert the correct name of your domain. The serverURL parameter indicates where clients should connect to for accessing our View Connection Server; replace the value of this parameter with the URI of your View Connection Server. Additional parameters are available in the client and can be set here. Refer to VMware View client documentation.

- g) Click the **Finished** button.
 - h) On the Variable Assign page, click the **Save** button.
22. On the Fallback path after **configureViewSSO**, click the **Deny** box.
23. Under Select Ending, click the **Allow** button, and then click **Save**.
24. On the Fallback path between **checkForViewClient** and **Deny**, click the + symbol.
25. In the General Purpose section, click the **Decision Box** option button, and then click the **Add Item** button. The Decision box page opens. Complete the following:
 - a) In the **Name** box, you can optionally type a name. In our example, we type **askUserDownload**.
 - b) In the **Message** box, type a message for users to see when the View client is not found. In our example, we type **View client not found**.
 - c) In the **Option 1** box, type something similar to **Download client now**.
 - d) In the **Option 2** box, type something similar to **Disconnect**.

- e) Click the **Save** button.
26. On the Option 1 path between **askUserDownload** and **Deny**, click the **+** symbol.
27. In the General Purpose section, click the **Resource Assign** option button, and then click **Add Item**. The Resource Assign page opens. Complete the following:
 - a) In the **Name** box, you can optionally type a new name. In our example, we type **DownloadViewClient**.
 - b) Click the **Add new entry** button.
 - c) Click the **Add/Delete Web Application Resources** link.
 - d) Check the box for the Web Application you created in *Creating a Web Application*, on page 1-8. In our example, we check the **DownloadViewClient** box. Click **Update**.
 - e) Click the **Save** button.
28. On the Fallback path after **DownloadViewClient**, click the **Deny** box.
29. Under Select Ending, click the **Allow** button, and then click **Save**.
30. 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.
31. Click the **Close** button on the upper right to close the VPE. Your policy should look similar to the following.

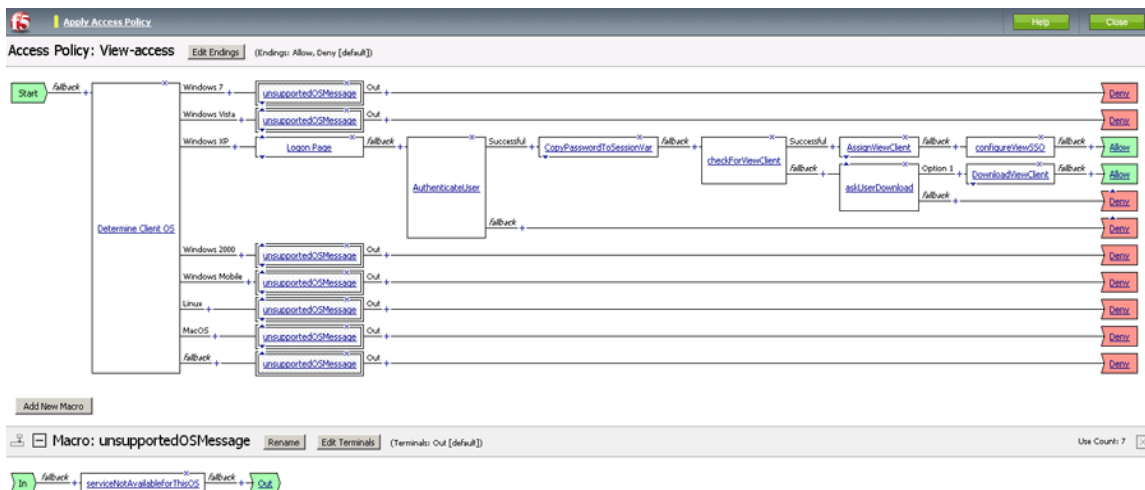


Figure 1.7 Completed Access Policy in the Visual Policy Editor

Creating the Network Access virtual server configuration objects

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. In the upper right portion of the screen, 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 **View-access-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 **View-access-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 **View-access-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

The next step is to create an SSL profile. This profile contains SSL certificate and Key information. 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 **View-access-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 iRule

The next task is to create an iRule that enables SSO for applications accepting user credentials via command line parameters.

To create the iRule

1. On the Main tab, expand **Local Traffic**, and then click **iRules**. The iRule screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New iRule screen opens.
3. In the Name box, type a name for this rule. In our example, we type **view-getUserLogonPassword**.
4. In the **Definition** box, copy and paste the iRule on the following page, omitting the line numbers.

```

1  when HTTP_REQUEST {
2      switch [HTTP::method] {
3          "POST" {
4              if { [HTTP::header Content-Type] eq "application/x-www-form-urlencoded" } {
5                  HTTP::collect [HTTP::header Content-Length]
6              }
7          }
8      }
9  }
10 when HTTP_REQUEST_DATA {
11     set namevals [split [HTTP::payload] "&"]
12     for {set i 0} {$i < [llength $namevals]} {incr i} {
13         foreach { name value } [split [lindex $namevals $i] "="] {
14             if { $name eq "password" } {
15                 set password [URI::decode $value]
16             }
17         }
18     }
19 }
20 when ACCESS_POLICY_AGENT_EVENT {
21     ACCESS::session data set session.vmware.sso $password
22 }
23 when ACCESS_POLICY_COMPLETED {
24     ACCESS::session data set session.vmware.sso ""
25 }

```

5. 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. In the upper right portion of the screen, 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 **View-tcp-443**.
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 1-15. We select **View-access-tcp-wan**. This is optional.
 9. From the **Protocol Profile (Server)** list, select the name of the profile you created in the *Creating the LAN optimized TCP profile* section. In our example, we select **View-access-tcp-lan**.
 10. From the **HTTP Profile** list, select the name of the profile you created in *Creating the HTTP profile*, on page 1-16. In our example, we select **View-access-http**.
 11. From the **SSL Profile (Client)** list, select the SSL profile you created in *Creating a Client SSL profile*, on page 1-16. In our example, we select **View-access-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 1-9. In our example, we select **View-access**.
 13. From the **Connectivity Profile** list, select the profile you created in *Creating a Connectivity Profile*, on page 1-6. In our example, we select **View-connectivity**.
 14. Leave the **Rewrite Profile** list set to **None**.
 15. **Do not** configure any of the options in the WAN Optimization section.
 16. In the **iRules** section, from the Available list, select the iRule you created in *Creating the iRule*, on page 1-17, and then click the Add (<<) button. In our example, we select **view-getUserLogonPassword**.
 17. Click the **Finished** button.
 18. 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**.
 - In Step 16, do **not** associate the iRule with this virtual server.
 - All other settings are the same.

This completes the BIG-IP Edge Gateway configuration.



2

Deploying the BIG-IP LTM with VMware View 4.5

Configuring the BIG-IP LTM for VMware View 4.5

In this section, we configure the BIG-IP LTM system for the VMware View 4.5 Connection Broker Servers.

Prerequisites and configuration notes

The following are prerequisites for this section:

- ◆ If you are using or plan to use PC over IP (PCoIP), see the *Special Note about PC over IP*, on page 2-1.
- ◆ Because the BIG-IP LTM is offloading SSL for the VMware deployment, this guide does not include VMware Security servers.
- ◆ This deployment guide is written with the assumption that VMware server(s), Virtual Center and connection brokers are already configured on the network and are in good working order.
- ◆ We recommend you enable direct connections to user's virtual desktops.
- ◆ For this deployment guide, the BIG-IP LTM system must be running version 10.2 or later. If you are using a previous version of the BIG-IP LTM system see the *Deployment Guide* index.
- ◆ **Important:** The current BIG-IP Application Template for View does not support PCoIP in View v4.0 or 4.5. We recommend using the following procedures for configuring the BIG-IP LTM with VMware View 4.0/4.5.

Special Note about PC over IP

Beginning with VMware View 4, VMware supports PC Over IP as a display protocol. PCoIP is an application encrypted UDP protocol, so the BIG-IP system cannot offload encryption for it.

If you want to use PCoIP, we recommend you enable direct connections to the desktop using PCoIP. This means the View client connects to the View Manager server for authentication, authorization and obtaining desktop information. Then, when the users choose a desktop to connect to, the view client opens a new connection directly to the desktop, bypassing the BIG-IP and connection manager. If you are deploying an environment with mixed display protocols, we recommend enabling direct access for all protocols. Refer to the VMware View administrators guide for details.

Modifying the VMware Virtual Desktop Manager global settings

Before starting the BIG-IP LTM configuration, we modify the View configuration to allow the BIG-IP LTM to load balance View client connections and offload SSL transactions. In the following procedure, we disable the SSL requirement for client connections in the Virtual Desktop Manager Administrator tool.

To modify the VMware configuration

1. Log on to the View Manager Administrator tool.
2. Click the **Configuration** tab.
The configuration options page opens.
3. In the Global Settings box, click the **Edit** button.
4. Clear the check from the **Require SSL for client connections** box.
5. Click the **OK** button.

◆ Note

This setting only applies to Connection Manager servers -- Security servers always require SSL

Configuring the External URL

The final modification to the VMware configuration is to configure the server External URL field with the FQDN of the BIG-IP virtual server. This is the server name that clients use to connect to the View Manager pool. Refer to the VMware View Administrator guide for more information. The following procedure must be performed on each VMware View Manager device.

To configure the External URL

1. Log on to the View Manager Administrator tool.
2. Click the **Configuration** tab.
3. Under **View Servers**, select a View Connection Server entry and click **Edit**.
4. In the **External URL** box, type the DNS name you will associate with the BIG-IP LTM virtual IP address, followed by a colon and the port. In our example, we type **https://broker.example.com:443**.
5. Clear the **Direct Connection to Desktop** box if it is checked.
6. Click the **Ok** button.

Configuring the BIG-IP LTM system for VMware Connection Brokers

In this section, we configure the BIG-IP LTM for the VMware View Connection Broker devices.

Creating the health monitor

The first task is to set up a health monitor for the VMware View Manager devices. This procedure is optional, but very strongly recommended. For this configuration, we create a simple HTTP health monitor. In this example, the advanced fields are not required, and we recommend you use the default values for the send and receive strings.

To configure a HTTP health monitor

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**. The Monitors screen opens.
2. Click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a name for the Monitor. In our example, we type **view-manager-http**.
4. From the **Type** list, select **HTTP**. The HTTP Monitor configuration options appear.
5. In the Configuration section, in the **Interval** and **Timeout** boxes, type an Interval and Timeout. We recommend at least a 1:3 +1 ratio between the interval and the timeout (for example, the default setting has an interval of **5** and an timeout of **16**). In our example, we use a **Interval** of **30** and a **Timeout** of **91**.
6. Click the **Finished** button. The new monitor is added to the Monitor list.

Creating the View Manager server pool

The next step is to create a pool on the BIG-IP LTM system for the View Manager systems. A BIG-IP pool is a set of devices grouped together to receive traffic according to a load balancing method.

To create the pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**. The Pool screen opens.
2. Click the **Create** button. The New Pool screen opens.
3. In the **Name** box, enter a name for your pool. In our example, we use **view-manager-pool**.

4. In the **Health Monitors** section, select the name of the monitor you created in *Creating the health monitor*, on page 2-3, and click the Add (<<) button. In our example, we select **view-manager-http**.
5. From the **Load Balancing Method** list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network). In our example, we select **Round Robin**.
6. For this pool, we leave the Priority Group Activation **Disabled**.
7. In the New Members section, make sure the **New Address** option button is selected.
8. In the **Address** box, add the View Manager to the pool. In our example, we type **10.133.80.10**
9. In the **Service Port** box, type **80**.
10. Click the **Add** button to add the member to the list.
11. Repeat steps 8-10 for each View Manager you want to add to the pool.
12. Click the **Finished** button.

Local Traffic » Pools » New Pool...

Configuration: Basic

Name: view-manager-pool

Health Monitors:

Active	Available
view-manager-http	gateway_icmp
	http
	https
	https_443
	inband

Resources:

Load Balancing Method: Round Robin

Priority Group Activation: Disabled

New Members:

New Address Node List

Address: 10.133.80.12

Service Port: 80 HTTP

Add

R:1 P:1 10.133.80.10 :80
R:1 P:1 10.133.80.11 :80
R:1 P:1 10.133.80.12 :80

Edit Delete

Cancel Repeat Finished

Figure 2.1 Configuring the BIG-IP LTM pool

Creating the Universal Inspection Engine persistence iRule

Using the following iRule, the BIG-IP LTM is able to direct traffic with greater precision resulting in a more uniform load distribution on the connection servers. Using the Universal Inspection Engine (UIE), the iRule looks for session information so that the BIG-IP LTM can persist the connections to the proper nodes. The View clients first use the session information in a cookie, and then use it as an URI argument when the tunnel is opened. The first response from the server contains a JSESSIONID cookie. The iRule enters that session ID into the connection table and upon further client requests looks for the information in a cookie or in the URI.

◆ Important

For the following iRule to function correctly, you must be using the BIG-IP LTM system to offload SSL transactions from the View implementation, as described in this deployment guide.

To create the persistence iRule

1. On the Main tab, expand **Local Traffic**, and then click **iRules**.
2. Click the **Create** button.
3. In the Name box, type a name for this rule. In our example, we type **view-jsessionid**.
4. In the Definition box, type the following iRule, omitting the line numbers.

Important: Line 19 must be entered as a single line.

```
1  when HTTP_REQUEST {
2      if { [HTTP::cookie exists "JSESSIONID"] } {
3          # log local0. "Client [IP::client_addr] sent cookie [HTTP::cookie "JSESSIONID"]"
4          set jsess_id [string range [HTTP::cookie "JSESSIONID"] 0 31]
5          persist uie $jsess_id
6          # log local0. "uie persist $jsess_id"
7      } else {
8          # log local0. "no JSESSIONID cookie, looking for tunnel ID"
9          set jsess [findstr [HTTP::uri] "tunnel?" 7]
10         if { $jsess != "" } {
11             # log local0. "uie persist for tunnel $jsess"
12             persist uie $jsess
13         }
14     }
15 }
16 when HTTP_RESPONSE {
17     if { [HTTP::cookie exists "JSESSIONID"] } {
18         persist add uie [HTTP::cookie "JSESSIONID"]
19         # log local0. "persist add uie [HTTP::cookie "JSESSIONID"] server: [IP::server_addr] client:
20         [IP::client_addr]"
21     }
22 }
23 # when LB_SELECTED {
24     # log local0. "Member [LB::server addr]"
25 }
```

5. Click the **Finished** button.

◆ **Tip**

The preceding iRule contains logging statements that are commented out. If you want to enable logging, simply remove the comment (#) from the code.

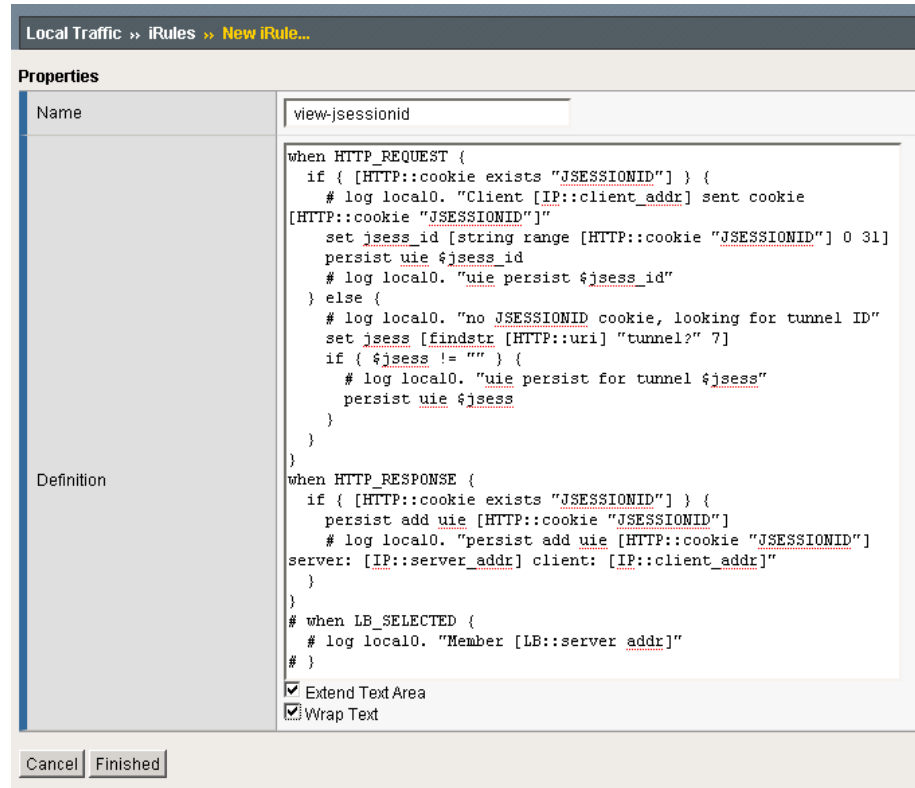


Figure 2.2 Configuring the persistence iRule on the BIG-IP LTM system

Using SSL certificates and keys

Before you can enable the BIG-IP LTM system to act as an SSL proxy, you must install a SSL certificate on the virtual server that you wish to use for client connections to the BIG-IP LTM device. For this deployment guide, we assume that you already have obtained an SSL certificate, but it is not yet installed on the BIG-IP LTM system. For information on generating certificates, or using the BIG-IP LTM to generate a request for a new certificate and key from a certificate authority, see the **Managing SSL Traffic** chapter in the *Configuration Guide for Local Traffic Management*.

Importing keys and certificates

Once you have obtained a certificate, you can import this certificate into the BIG-IP LTM system using the Configuration utility. By importing a certificate or archive into the Configuration utility, you ease the task of managing that certificate or archive. You can use the Import SSL Certificates and Keys screen only when the certificate you are importing is in Privacy Enhanced Mail (PEM) format.

To import a key or certificate

1. On the Main tab, expand **Local Traffic**.
2. Click **SSL Certificates**. The list of existing certificates displays.
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.

Creating BIG-IP LTM profiles

The next task is to create the profiles. A *profile* is an object that contains user-configurable settings for controlling the behavior of a particular type of network traffic, such as HTTP connections. Using profiles enhances your control over managing network traffic, and makes traffic-management tasks easier and more efficient.

Although it is possible to use the default profiles, we strongly recommend you create new profiles based on the default parent profiles, even if you do not change any of the settings initially. Creating new profiles allows you to easily modify the profile settings specific to this deployment, and ensures you do not accidentally overwrite the default profile.

Creating an HTTP profile

The first new profile we create is an HTTP profile. The HTTP profile contains numerous configuration options for how the BIG-IP LTM system handles HTTP traffic. In this example, we use the **http-lan-optimized-caching** parent profile.

To create a new HTTP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.

2. In the upper right portion of the screen, click the **Create** button. The New HTTP Profile screen opens.
3. In the **Name** box, type a name for this profile. In our example, we type **view-http**.
4. From the **Parent Profile** list, select **http-lan-optimized-caching**. The profile settings appear.
5. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
6. Click the **Finished** button.

Creating the TCP profiles

The next task is to create the TCP profiles. We recommend creating one TCP profile using the **tcp-lan-optimized** parent. If your configuration uses various WAN links and your users are widely distributed, you should also create a second profile that uses **tcp-wan-optimized** as the parent profile. If all of your users are accessing the BIG-IP LTM over a LAN, you only need to create the LAN optimized profile.

Creating the LAN optimized TCP profile

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

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Protocol** menu, click **tcp**.
3. In the upper right portion of the screen, click the **Create** button.
4. In the **Name** box, type a name for this profile. In our example, we type **view-lan-opt**.
5. From the **Parent Profile** list, select **tcp-lan-optimized**.
6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating the WAN optimized TCP profile

Now we create is the WAN optimized TCP profile.

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. On the Menu bar, from the **Protocol** menu, click **tcp**.
3. In the upper right portion of the screen, click the **Create** button.

-
4. In the **Name** box, type a name for this profile. In our example, we type **view-wan-opt**.
 5. From the **Parent Profile** list, select **tcp-wan-optimized**.
 6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
 7. Click the **Finished** button.

Creating the UIE persistence profile

The next profile we create is the persistence profile. This profile references the Universal Inspection Engine iRule you created earlier in this guide.

To create a persistence profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, click **Persistence**. The Persistence Profiles screen opens.
3. In the upper right portion of the screen, click the **Create** button. The New Persistence Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **view-persist**.
5. From the **Persistence Type** list, select **Universal**. The configuration options for universal persistence appear.
6. In the **iRule** row, check the Custom box. From the iRule list, select the name of the iRule you created in *Creating the Universal Inspection Engine persistence iRule*, on page 2-5. In our example, we select **view-jsessionid**.
7. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
8. Click the **Finished** button.

Figure 2.3 Creating the persistence profile

Creating a OneConnect profile

The next profile we create is a OneConnect profile. With OneConnect enabled, client requests can use existing, server-side connections, thus reducing the number of server-side connections that a server must open to service those requests. For more information on OneConnect, see the BIG-IP LTM documentation.

In our example, we leave all the options at their default settings. You can configure these options as appropriate for your network.

To create a new OneConnect profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Other** menu, click **OneConnect**. The Persistence Profiles screen opens.
3. In the upper right portion of the screen, click the **Create** button. The New HTTP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **view-oneconnect**.
5. From the **Parent Profile** list, ensure that **oneconnect** is selected.
6. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Creating a Client SSL profile

The next step in this configuration is to create a Client SSL profile. This profile contains the SSL certificate and Key information for offloading the SSL traffic.

To create a new Client SSL profile based on the default profile

1. On the Main tab, expand **Local Traffic**.
2. Click **Profiles**.
The HTTP Profiles screen opens.
3. On the Menu bar, from the SSL menu, select **Client**.
The Client SSL Profiles screen opens.
4. In the upper right portion of the screen, click the **Create** button.
The New Client SSL Profile screen opens.
5. In the **Name** box, type a name for this profile. In our example, we type **view-clientssl**.
6. In the Configuration section, check the **Certificate** and **Key Custom** boxes.
7. From the **Certificate** list, select the name of the Certificate you imported in the *Importing keys and certificates* section.
8. From the **Key** list, select the key you imported in the *Importing keys and certificates* section.
9. Click the **Finished** button.

Creating the virtual server

Next, we configure a virtual server that references the profiles and pool you created in the preceding procedures.

To create the virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
The Virtual Servers screen opens.
2. In the upper right portion of the screen, 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 **view-virtual**.
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 **10.133.81.10**.
6. In the **Service Port** box, type **443**, or select **HTTPS** from the list.

General Properties	
Name	view-virtual
Destination	Type: <input checked="" type="radio"/> Host <input type="radio"/> Network Address: 10.133.81.10
Service Port	443 HTTPS
State	Enabled

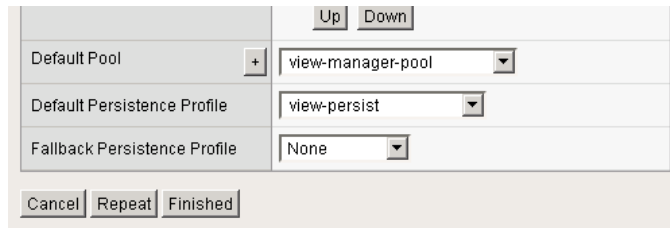
Figure 2.4 Configuring the virtual server properties

7. From the Configuration list, select **Advanced**.
The Advanced configuration options appear.
8. Leave the **Type** list at the default setting: **Standard**.
9. From the **Protocol Profile (Client)** list select the profile you created in *Creating the WAN optimized TCP profile*, on page 2-8. In our example, we select **view-wan-opt**.
10. From the **Protocol Profile (Server)** list, select the profile you created in *Creating the UIE persistence profile*, on page 2-9. In our example, we select **view-lan-opt**.
11. From the **OneConnect Profile** list, select the profile you created in *Creating a OneConnect profile*, on page 2-10. In our example, we select **view-oneconnect**.
12. From the **HTTP Profile** list, select the profile you created in *Creating an HTTP profile*, on page 2-7. In our example, we select **view-http**.

Configuration: Advanced	
Type	Standard
Protocol	TCP
Protocol Profile (Client)	view-wan-opt
Protocol Profile (Server)	view-lan-opt
OneConnect Profile	view-oneconnect
NTLM Conn Pool	None
HTTP Profile	view-http
FTP Profile	None
SSL Profile (Client)	view-clientssl
SSL Profile (Server)	view-serverssl

Figure 2.5 Adding the profiles to the virtual server

-
13. In the Resources section, from the **Default Pool** list, select the pool you created in *Creating the View Manager server pool*, on page 2-3. In our example, we select **view-manager-pool**.
 14. From the **Default Persistence Profile** list, select the persistence profile you created in *Creating the UIE persistence profile*, on page 2-9. In our example, we select **view-persist**.



		Up	Down
Default Pool	+	view-manager-pool	
Default Persistence Profile		view-persist	
Fallback Persistence Profile		None	
Cancel		Repeat	Finished

Figure 2.6 Adding the pool and persistence profile to the virtual server

15. Click the **Finished** button.

The BIG-IP LTM configuration for VMware View is now complete.