



# Deploying F5 with Microsoft Forefront Unified Access Gateway



**Microsoft**® Partner

---

# Introducing the F5 and Microsoft Unified Access Gateway configuration

Welcome to the F5 and Microsoft® Forefront™ Unified Access Gateway (UAG) deployment guide. This guide contains step-by-step procedures for configuring F5 devices for Forefront UAG resulting in a secure, fast and highly available deployment.

Forefront Unified Access Gateway (UAG) provides remote access to corporate resources from a diverse range of remote client endpoints and locations. By using Forefront UAG, you can publish Web and non-Web applications, and provide access to internal networks. You can control and help secure endpoint access with a number of control mechanisms, including client authentication, application authorization, and endpoint health validation against access policies.

Following installation, Forefront UAG servers can be configured as Forefront UAG DirectAccess servers, providing remote users with the experience of being seamlessly connected to your internal networks, extending the benefits of Windows DirectAccess across your infrastructure, enhancing scalability, and simplifying deployments and ongoing management.

For more information on Microsoft Forefront Unified Access Gateway, see <http://www.microsoft.com/forefront/prodinfo/roadmap/uag.mspx>

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

You can also visit the Microsoft page of F5's online developer community, DevCentral, for Microsoft forums, solutions, blogs and more:

<http://devcentral.f5.com/Default.aspx?tabid=89>.

## Prerequisites and configuration notes

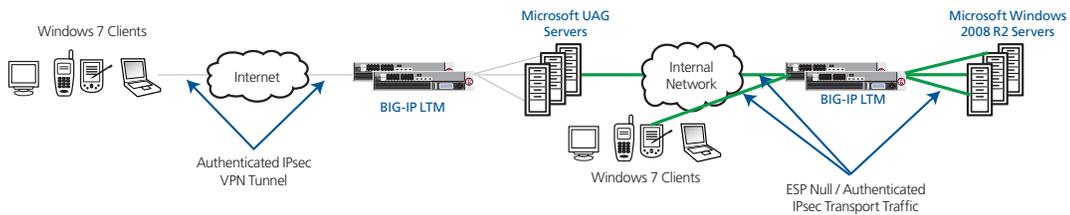
The following are general prerequisites for this deployment.

- ◆ All of the configuration procedures in this document are performed on the BIG-IP Local Traffic Manager (LTM) system. For information on how to deploy or configure the Unified Access Gateway, consult the appropriate Microsoft documentation.
- ◆ This document is written with the assumption that you are familiar with both the F5 devices and Microsoft Forefront Unified Access Gateway. For more information on configuring these products, consult the appropriate documentation.
- ◆ The BIG-IP LTM system should be running version 9.4 or later. This guide was tested with version 10.0.1.
- ◆ The following BIG-IP LTM configuration instructions assume you are connected to the web-based configuration utility using a web browser.

- ◆ The F5 defines the term *connection* as a flow of packets between a source and destination. By this definition, UDP traffic is considered a connection. This differs from the traditional (and still accurate) definition of UDP which is considered *connectionless*.

## Configuration example

F5's BIG-IP LTM can be used to provide scalability and high availability for Microsoft's Unified Access Gateway. When deployed on either side of the UAG servers, BIG-IP's load balancing capabilities can be leveraged to route both incoming and outgoing traffic through the most appropriate UAG server.



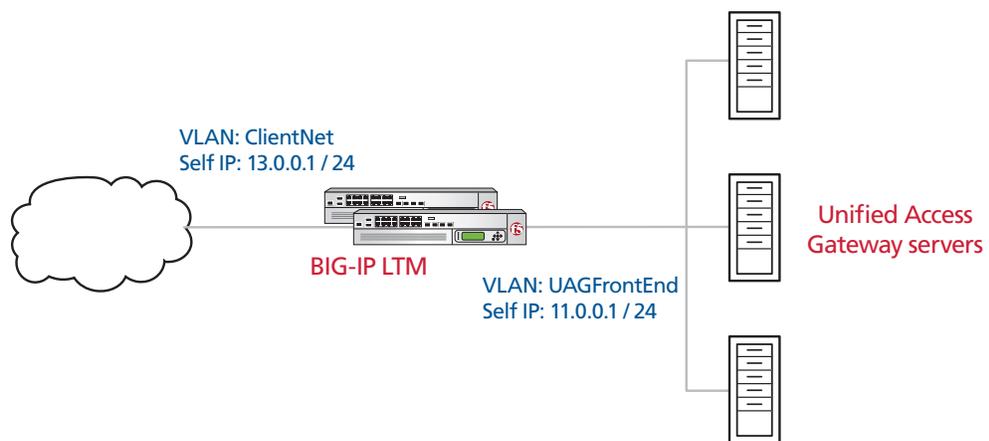
**Figure 1** Logical configuration example

A unique requirement of a scaled UAG deployment is that server generated connections to an external Client are routed to the UAG server where the Client has a pre-established tunnel. Without doing this, these server generated connections could be sent to the wrong UAG server and the connection would be dropped. The BIG-IP LTM handles this traffic with its intelligent traffic engine, iRules, to track client-to-UAG server tunnels, and match server generated connections to the right UAG server.

---

## Configuring the external BIG-IP LTM system

Both the external and internal BIG-IP systems manage bi-directional traffic flows. The external BIG-IP system is configured to accept incoming connections from the clients and load balance them to the available UAG servers. By creating a forwarding Virtual Server, the BIG-IP system also forwards connections originating from the servers that are destined to systems on the external network.



*Figure 2 External BIG-IP LTM logical configuration*

## Network configuration of the external BIG-IP

The external BIG-IP needs two VLANs, one facing the external client network and the other connected to the UAG front ends. In this example, we have named our VLANs as follows.

- **ClientNet**  
This VLAN faces the external cloud, and is the network that provides access to and from the clients. In most cases, this is the Internet VLAN.
- **UAGFrontEnd**  
The VLAN connects the BIG-IP to the front end of the UAG servers.

Each VLAN on the BIG-IP also needs an administrative IP address and netmask. These IPs are considered 'Self IP' addresses, and are used for a handful of administrative and routing tasks.

We used the following IP addresses for testing our UAG/BIG-IP solution; yours will most likely vary:

- **ClientNet** - 13.0.0.1 /24
- **UAGFrontEnd** - 11.0.0.1 /24

**Configuration Tasks:**

- *Creating the VLANs*
- *Creating the Self IP addresses*

## Creating the VLANs

In this section, we create both the ClientNet and UAGFrontEnd VLANs on the BIG-IP LTM system.

**To create the ClientNet VLAN**

1. On the Main tab, expand **Network**, and then click **VLANs**. The VLAN screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New VLAN screen opens.
3. In the **Name** box, enter a name for your VLAN. In our example, we use **ClientNet**.
4. In the **Interface** section, from the **Available** list, click the physical interface associated with this VLAN and click the Add button (<< or >>) to move it to the **Untagged** or **Tagged** box. Repeat this step as necessary.

In our example, port **1.1** was connected to the ClientNet VLAN. In your deployment, this may differ.

The screenshot shows the 'New VLAN...' configuration window. The breadcrumb path is 'Network >> VLANs >> New VLAN...'. Under 'General Properties', the 'Name' field contains 'ClientNet' and the 'Tag' field is empty. Under 'Resources', the 'Interfaces' section shows three columns: 'Untagged', 'Available', and 'Tagged'. The 'Available' list contains interfaces 1.1 through 1.6, with 1.1 selected and moved to the 'Untagged' column. Below this, the 'Configuration' section is set to 'Basic', with 'Source Check' unchecked and 'MTU' set to 1500. At the bottom are 'Cancel', 'Repeat', and 'Finished' buttons.

*Figure 3* Creating the ClientNet VLAN

5. Click the **Repeat** button.

6. Repeat steps 3 and 4. In step 3, give the VLAN a unique name. In our example, we use **UAGFrontEnd**. In step 4, use the appropriate interface. In our example, we use **1.2**.
7. Click the **Finished** button.

When you have finished creating the two VLANs, your VLAN list should look similar to Figure 4.

<input type="checkbox"/>	Name	Tag	Untagged Interfaces	Tagged Interfaces
<input type="checkbox"/>	ClientNet	4094	1.1	
<input type="checkbox"/>	UAGFrontEnd	4093	1.2	

*Figure 4 The VLAN List*

## Creating the Self IP addresses

The next task is to create a Self IP address on the BIG-IP LTM for each of the VLANs you just created.

### To create a Self IP address

1. On the Main tab, expand **Network**, and then click **Self IPs**.
2. Click the **Create** button. The New Self IP screen opens.
3. In the **IP Address** box, type an IP Address for the first VLAN you created. In our example, we use **13.0.0.1** for the ClientNet VLAN.
4. In the **Netmask** box, enter a netmask for the Self IP. In our example, we use **255.255.255.0**
5. From the **VLAN** list, select the name of the first VLAN you created. In our example, we select **ClientNet**.

Configuration	
IP Address	13.0.0.1
Netmask	255.255.255.0
VLAN	ClientNet
Port Lockdown	Allow Default

Buttons: Cancel, Repeat, Finished

*Figure 5 Creating a Self IP address*

6. Click the **Repeat** button.
7. Return to step 2 and repeat this procedure for the second VLAN you created, using the appropriate IP address and Netmask, and selecting the 2nd VLAN you created from the VLAN list. In our example, we use **11.0.0.1** for the IP address, **255.255.255.0** for the Netmask, and we select **UAGFrontEnd** from the VLAN list.
8. Click the **Finished** button.

When you have finished creating the two Self IPs, the Self IP list should look like Figure 6.

<input checked="" type="checkbox"/>	IP Address	Netmask	VLANs	Unit ID
<input type="checkbox"/>	11.0.0.1	255.255.255.0	UAGFrontEnd	0
<input type="checkbox"/>	13.0.0.1	255.255.255.0	ClientNet	0

*Figure 6 Self IP list*

## Configuring Load Balancing for Incoming Connections

The external BIG-IP needs to be configured to accept incoming connections, regardless of the protocol, and load balance them across the UAG farm. By configuring the BIG-IP to monitor the health and availability of the UAG servers, the BIG-IP is able to efficiently send connections to the most appropriate UAG server.

In this section, we configure the BIG-IP to accept and load balance connections using Teredo and IP-HTTPS to the appropriate UAG servers.

### Configuration Tasks

- *Creating the UAG Teredo Monitor*
- *Creating the Teredo pools*
- *Creating the virtual servers*

## Creating the UAG Teredo Monitor

The most effective way to monitor the UAG front ends for Teredo availability is by using a combination of monitors to make sure the Teredo service is accepting connections. We recommend the default UDP monitor checking port 3544 and the default HTTPS monitor checking port 443.

---

Although the default interval and timeout values (5 and 16) for the monitors provides the optimal behavior, these values can be modified for more aggressive or relaxed monitoring. In this guide, we recommend using the default settings.

Because we recommend using the default UDP and HTTPS monitors, there are no step-by-step procedures. If you would like to use a custom monitor to change specific settings, refer to the BIG-IP administrator's guide.

## Creating the Teredo pools

Each of the UAG front ends has two unique and sequential IPv4 addresses for the BIG-IP LTM to send the Teredo traffic. The primary Teredo IP on each UAG server is the initial and preferred for traffic, and the secondary IP is the backup.

Because of this design, we create two pools on the BIG-IP for the Teredo traffic. The first pool is a collection of the primary IP addresses, and the second pool is a collection of the secondary IPs.

### To create the Teredo pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**. The Pool screen opens.
2. In the upper right portion of the screen, 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 **UAG\_TEREDO\_PRIMARY**.
4. In the **Health Monitors** section, from the Available list, select **UDP**, and click the Add (<<) button. Next, select the **HTTPS\_443** monitor, and click the Add (<<) button.
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). We recommend you use a dynamic load balancing method (least connections, observed, predictive) in order to have BIG-IP factor in server and network performance in its load balancing decision.
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 first server to the pool. In our example, we type **11.0.0.40**.
9. In the Service Port box, type **3544**.
10. Click the **Add** button to add the member to the list.
11. Repeat steps 8-10 for each server you want to add to the pool.

**Figure 7** Creating the primary Teredo pool

12. Click the **Repeat** button. As the Teredo services on each UAG server also have a secondary IP, the BIG-IP must be configured with a pool of these IPs.
13. Return to Step 3 and repeat this procedure for the secondary Teredo IP. Use an appropriate name and the relevant IP addresses. The port number and health monitors are the same.

◆ **Note**

*The BIG-IP system supports a wide range of load balancing methods to suit a variety of traffic types. By design, Teredo connections span a wide range of life times, which can make static load balancing methods such as Round Robin less desirable. The BIG-IP includes load balancing methods built to suit connection behavior like this, and they are built upon the **least connections** model. It is recommended that for best performance, either Least Connections, Observed, or Predictive load balancing methods are used for Teredo traffic. More information on the available load balancing methods can be found in the BIG-IP Configuration Guide for Local Traffic Management.*

---

## Creating the virtual servers

After the two Teredo pools have been created, we create primary and secondary virtual servers (VIP) for Teredo connections. These virtual IP addresses then point to their respective primary and secondary Teredo pools. Just like the Teredo servers require two sequential IPs for access, we want the BIG-IP to have two sequential IP addresses for the primary and secondary VIPs. The Virtual IPs listen on port 3544 for UDP traffic.

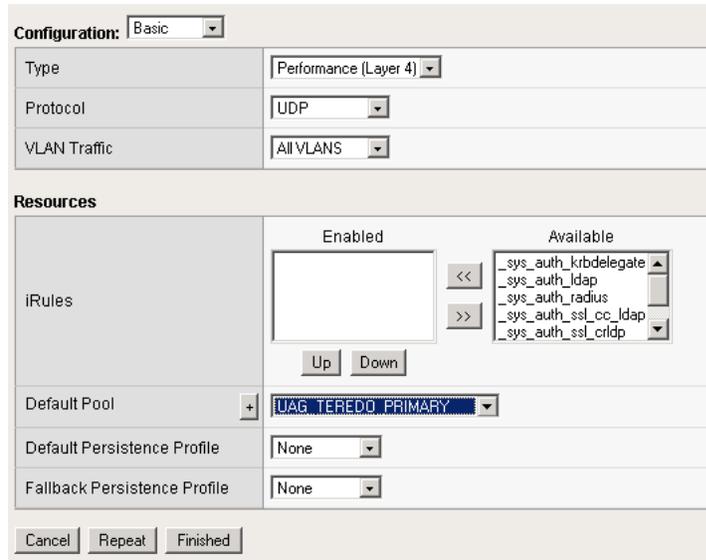
### To create the Teredo virtual servers

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 **TEREDO\_PRIMARY\_VIP**.
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 **13.0.0.30**.
6. In the Service Port box, type **3544**.

General Properties	
Name	TEREDO_PRIMARY_VIP
Destination	Type: <input checked="" type="radio"/> Host <input type="radio"/> Network Address: 13.0.0.30
Service Port	3544 <input type="button" value="Select..."/>
State	Enabled <input type="button" value="v"/>

*Figure 8 The general properties of the virtual server*

7. In the Configuration section, from the **Type** list, select **Performance (Layer 4)**.
8. From the Protocol list, select **UDP**.
9. From the Default Pool list, select the name of your primary pool. In our example, we select **UAG\_TEREDO\_PRIMARY**.



**Figure 9** Configuration and Resources sections of the virtual server

10. Click the **Repeat** button to create the secondary Teredo virtual server.
11. Return to Step 3 and repeat this procedure for the secondary Teredo virtual server. Use an appropriate name, relevant IP address, and select the second Teredo pool you created. The rest of the configuration options should remain the same.
12. Click the **Finished** button.

Now you should have two sequential VIPs for Teredo connections, and your virtual server list should look like Figure 10.



**Figure 10** the virtual server list showing the Teredo virtual servers

---

## Configuring the BIG-IP for access via IP-HTTPS

The external BIG-IP system can also be configured to allow access from clients using IP-HTTPS as their connection protocol. In this case, a BIG-IP Virtual IP listening on port 443 load balances connections to the set of IP-HTTPS enabled UAG front ends.

### Configuration Tasks:

- *Creating the UAG IP-HTTPS monitor*
- *Creating the pool of IP-HTTPS servers*
- *Creating the IP-HTTPS virtual server*

## Creating the UAG IP-HTTPS monitor

IP-HTTPS monitoring of the UAG front ends can be done at the application layer by using the default HTTPS monitor on the BIG-IP. This ensures a valid TCP handshake, SSL negotiation, and basic data request/response can be made before the BIG-IP assumes the UAG front end is ready for IP-HTTPS traffic.

Although the default monitor is recommended, it is possible to modify the interval and timeout values in order to configure the monitor to be more aggressive or relaxed.

Since use of the default HTTPS monitor is being recommended, no steps for creating the monitor are being shown.

## Creating the pool of IP-HTTPS servers

For IP-HTTPS load balancing, first create the pool populated with the IPs of the IP-HTTPS enabled UAG servers.

1. On the Main tab, expand **Local Traffic**, and then click **Pools**. The Pool screen opens.
2. In the upper right portion of the screen, 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 **UAG\_IPHTTPS\_POOL**.
4. In the **Health Monitors** section, from the **Available** list, select the **HTTPS** monitor, and click the Add (<<) button.
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). We recommend you use a dynamic load balancing method (least connections, observed, predictive) in order to have BIG-IP factor in server and network performance in its load balancing decision.
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 first server to the pool. In our example, we type **11.0.0.45**.
9. In the **Service Port** box, type **443**.
10. Click the **Add** button to add the member to the list.
11. Repeat steps 8-10 for each server you want to add to the pool.
12. Click the **Finished** button.

*Figure 11 Creating the HTTPS pool*

## Creating the IP-HTTPS virtual server

For IP-HTTPS traffic, we configure the virtual server to listen on port 443 for incoming connections. The virtual server should be configured to use TCP as the protocol, and reference the IP-HTTPS pool you just created.

### To create the Teredo virtual servers

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 **IPHTTPS\_VIP**.
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 **13.0.0.35**.
6. In the Service Port box, type **443** or select HTTPS from the list.
7. In the Configuration section, from the **Type** list, select **Performance (Layer 4)**.
8. From the **Protocol** list, select **TCP**.
9. From the **Default Pool** list, select the name of the pool you just created. In our example, we select **UAG\_IPHTTPS\_POOL**.
10. Click the **Finished** button.

The screenshot shows a configuration window for a virtual server. At the top, there's a 'Configuration:' dropdown set to 'Basic'. Below it are several rows of configuration options: 'Type' is 'Performance (Layer 4)', 'Protocol' is 'TCP', and 'VLAN Traffic' is 'All VLANs'. A 'Resources' section follows, containing an 'iRules' table with 'Enabled' and 'Available' columns. The 'Available' column lists several system authentication rules. Below the iRules table are 'Up' and 'Down' buttons. Further down, 'Default Pool' is set to 'UAG\_IPHTTPS\_POOL', and both 'Default Persistence Profile' and 'Fallback Persistence Profile' are set to 'None'. At the bottom, there are three buttons: 'Cancel', 'Repeat', and 'Finished'.

*Figure 12* Configuring the IP-HTTPS virtual server

You should now have three virtual servers (two for Teredo and one for IP-HTTPS).

## Configuring the BIG-IP LTM to forward outbound connections

Connections arriving from the UAG servers to the external BIG-IP need to be forwarded onto their destination. To enable this forwarding, we create a network forwarding Virtual IP on the BIG-IP.

### To create the forwarding 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 **FORWARDING\_TO\_CLIENTS**.
4. In the Destination section, select the **Network** option button.
5. In the **Address** box, type **0.0.0.0**.
6. In the **Mask** box, type **0.0.0.0**.
7. In the **Service Port** box, type **0** or select **\*All Ports** from the list.

The screenshot shows the 'New Virtual Server' configuration window. Under 'General Properties', the 'Name' field contains 'FORWARDING\_TO\_CLIENTS'. The 'Destination' section has 'Type' set to 'Network' (selected with a radio button), 'Address' set to '0.0.0.0', and 'Mask' set to '0.0.0.0'. The 'Service Port' section has a dropdown menu set to '\*All Ports'. The 'State' dropdown is set to 'Enabled'.

*Figure 13* Configuring the forwarding virtual server

8. In the Configuration section, from the **Type** list, select **Forwarding (IP)**.
9. From the **Protocol** list, select **\*All Protocols**.
10. Click the **Finished** button.

The screenshot shows the 'Resources' configuration section. Under 'Configuration', 'Type' is 'Forwarding (IP)', 'Protocol' is '\*All Protocols', and 'VLAN Traffic' is 'All VLANs'. Under 'Resources', there are two columns: 'Enabled' (empty) and 'Available'. The 'Available' column contains a list of iRules: '\_sys\_auth\_krbdelegate', '\_sys\_auth\_ldap', '\_sys\_auth\_radius', '\_sys\_auth\_ssl\_cc\_ldap', and '\_sys\_auth\_ssl\_crdp'. Navigation buttons '<<', '>>', 'Up', and 'Down' are visible between the columns. At the bottom, there are 'Cancel', 'Repeat', and 'Finished' buttons.

*Figure 14* Configuring the forwarding virtual server resources

As a security measure, we want to enable this VIP only on the internal VLAN. It should remain disabled on the external VLAN. The final Virtual IP list should look like Figure 15.

Local Traffic » Virtual Servers

Virtual Server List Virtual Address List Statistics

Search

<input type="checkbox"/>	Status	Name	Partition	Destination	Service Port	Type
<input type="checkbox"/>	<span style="color: blue;">■</span>	FORWARDING_TO_CLIENTS	Common	0.0.0.0	0 (Any)	Forwarding (IP)
<input type="checkbox"/>	<span style="color: green;">●</span>	IP-HTTPS	Common	10.0.0.32	443 (HTTPS)	Standard
<input type="checkbox"/>	<span style="color: green;">●</span>	TEREDO_PRIMARY	Common	10.0.0.30	3544	Standard
<input type="checkbox"/>	<span style="color: green;">●</span>	TEREDO_SECONDARY	Common	10.0.0.31	3544	Standard

Enable Disable Delete...

*Figure 15 Final virtual server list for the external BIG-IP system*

## Configuring the internal BIG-IP LTM system

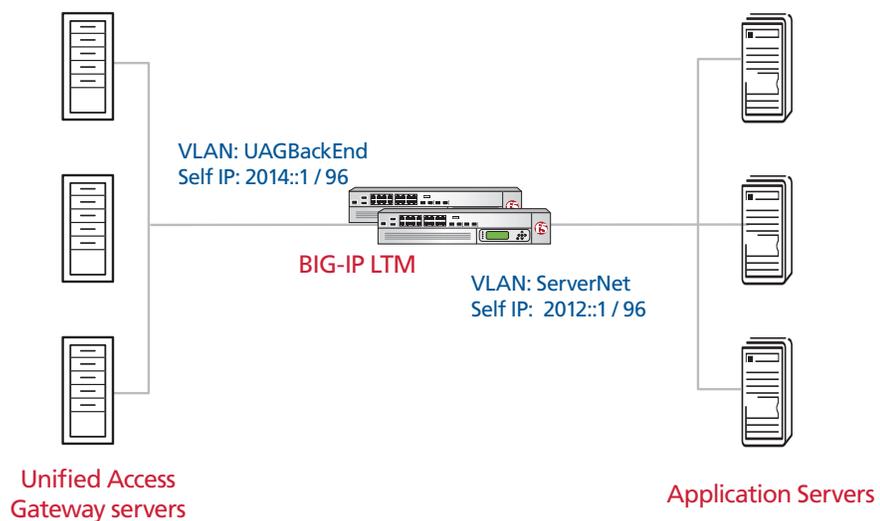
The internal BIG-IP is configured to pass client to server traffic onto the destined server. By leveraging BIG-IP's Auto Lasthop functionality, return packets from the servers are sent to the same UAG server to which the client was originally sent.

We also use iRules to make sure new server-generated connections to an external system are sent to the same UAG server to which the system already has an established tunnel.

## Network configuration of the internal BIG-IP

The Internal BIG-IP also uses two VLANs. One VLAN connects to the internal side of the UAG servers, and the other VLAN connects to the internal server network. In this example, we have named our VLANs

- **UAGBackEnd** - This VLAN connects the BIG-IP to the UAG servers.
- **ServerNet** - This VLAN connects the BIG-IP to the servers.



**Figure 16** Internal BIG-IP LTM system logical configuration example

You need to configure the BIG-IP with an administrative IP address and netmask for each VLAN. These IPs are considered 'Self IP' addresses, and are used for a handful of administrative and routing tasks. The IP addresses in the following example represent what we used for the testing of this solution, yours will most likely vary.

- **UAGBackEnd**  
2014.0.0.0.0.0.1 /96  
14.0.0.1 /24
- **ServerNet**  
2012.0.0.0.0.0.1 /96  
12.0.0.1 /24

**Configuration Tasks:**

- *Creating the VLANs*
- *Creating the Self IP addresses*

## Creating the VLANs

To create the VLANs, follow the procedure *Creating the VLANs*, on page 4. Use an easily identifiable name, and an appropriate physical interface (see the previous section for our examples). Repeat this procedure twice, once for the back end VLAN and once for the internal server network VLAN. In our example, we create VLANs named **ServerNet** and **UAGBackEnd**.

## Creating the Self IP addresses

To create the Self IP addresses, follow the procedure *Creating the Self IP addresses*, on page 5. Use an appropriate IP address and netmask, and select the corresponding VLAN from the list. Create a Self IP address for each of the VLANs you created in the preceding procedure. In our example, our UAGBackEnd Self IP has an IP address of **2014::1** and a netmask of **ffff.ffff.ffff.ffff::** (see Figure 17). The ServerNet VLAN Self IP has a IP address of 2012::1 and a netmask of **ffff.ffff.ffff.ffff::**.

The screenshot shows a configuration window titled "Network >> Self IPs". Under the "Configuration" section, there are four rows of fields:

IP Address	2014::1
Netmask	ffff.ffff.ffff.ffff::
VLAN	UAGBackEnd
Port Lockdown	Allow Default

At the bottom of the configuration area, there are three buttons: "Cancel", "Repeat", and "Finished".

**Figure 17** *Configuring the IPv6 backend Self IP address*

## Forwarding server-bound connections

Client generated connections that are bound for the application servers will have already been load balanced by the external BIG-IP, and it is the responsibility of the internal BIG-IP to track which UAG server the packets came from, and forward them on to the application servers. The tracking is done by an iRule, and the forwarding is done by a forwarding Virtual IP.

### Configuration Tasks:

- *Creating the connection tracking iRule*
- *Creating the forwarding IPv4 virtual server*
- *Creating the IPv6 forwarding virtual server*

## Creating the connection tracking iRule

The connection tracking iRule records the source IP address of the connection and the MAC address of the sending UAG server into an internal table on the BIG-IP. This information is used later to make sure that a connection to the same client is always sent to the same UAG server.

### 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, enter a name for your iRule. In our example, we use **linkpersist\_add**.
4. In the Definition section, copy and paste the following iRule:

```
when SERVER_DATA {  
    set session_key [IP::local_addr]  
    session add uie {$session_key any virtual} [LINK::nexthop] 43200  
    log local0. "fletching, session table entry added for $session_key to [LINK::nexthop]"  
}
```

5. Click **Finished**.

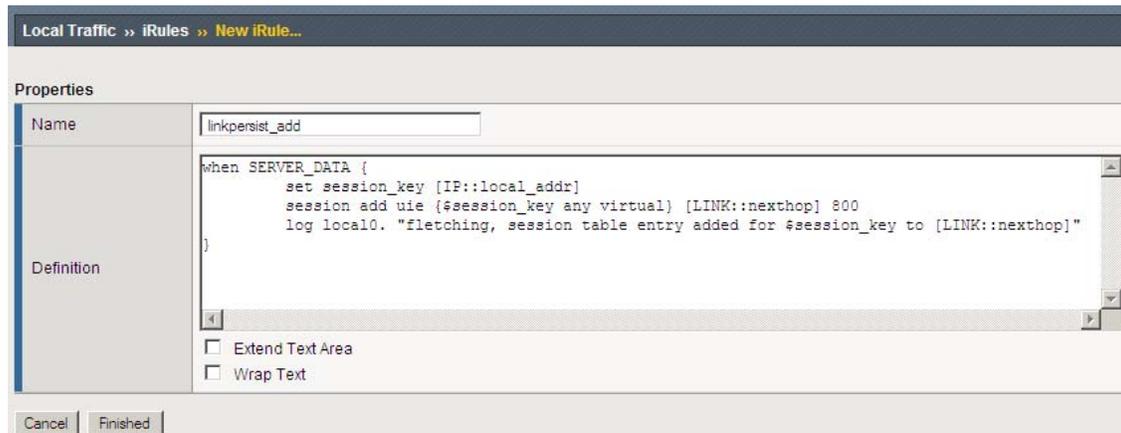


Figure 18 Creating the Forwarding iRule

## Creating the forwarding IPv4 virtual server

The IPv4 forwarding VIP forwards all IPv4 connections.

### To create the forwarding virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this virtual server. In our example, we type **ForwardingIPv4**.
4. In the Destination section, select the **Network** option button.
5. In the **Address** box, type **0.0.0.0**.
6. In the **Mask** box, type **0.0.0.0**.
7. In the **Service Port** box, type **0** or select All Ports
8. In the Configuration section, from the **Type** list, select **Forwarding (IP)**.
9. From the **Protocol** list, select **\*All Protocols**.
10. From the **VLAN Traffic** list, select **Enabled On...** The VLAN List row appears.
11. From the VLAN List **Available** list, select the VLAN you created for the back end in *Creating the VLANs*, on page 17, and click the Add (<<) button. In our example, we select **UAGBackEnd**.
12. In the Resources section, from the iRules **Available** list, select the name of the iRule you created in *Creating the connection tracking iRule*, on page 18, and click the Add (<<) button. In our example, we select **linkpersist\_add**.
13. Click **Finished**.

## Creating the IPv6 forwarding virtual server

The next task is to create an IPv6 forwarding virtual server.

### To create the IPv6 virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
2. In the upper right portion of the screen, click the **Create** button.
3. In the **Name** box, type a name for this virtual server. In our example, we type **ForwardingIPv6**.
4. In the Destination section, select the **Network** option button.
5. In the Address box, type **0.0.0.0.0.0.0**.
6. In the Mask box, type **0.0.0.0.0.0.0**.
7. In the **Service Port** box, type **0**.
8. In the Configuration section, from the Type list, select **Forwarding (IP)**.
9. From the Protocol list, select **\*All Protocols**.
10. From the **VLAN Traffic** list, select **Enabled On...** The VLAN List row appears.
11. From the VLAN List **Available** list, select the VLAN you created for the back end in *Creating the VLANs*, on page 17, and click the Add (<<) button. In our example, we select **UAGBackEnd**.
12. In the Resources section, from the iRules **Available** list, select the name of the iRule you created in *Creating the connection tracking iRule*, on page 18, and click the Add (<<) button. In our example, we select **linkpersist\_add**.
13. Click **Finished**.

## Load balancing client-bound connections

Connections generated by a member of the internal server network to an external system need to be sent to the same UAG server that the external system has a previously established tunnel. The iRule created above, along with a new iRule here, ensures that the correct UAG server is selected.

### Configuration Tasks:

- *Creating the UAG Server Monitor*
- *Creating the pool of UAG Servers (IPv6)*
- *Creating the pre-selection iRule*
- *Creating the virtual server*

---

## Creating the UAG Server Monitor

In order to load balance connections to the UAG servers, the BIG-IP system first check device availability. In this case, we use the default ICMP monitor. If you want to create a new monitor, we recommend you use the ICMP parent monitor.

## Creating the pool of UAG Servers (IPv6)

The next step is to create a pool on the BIG-IP system for the UAG servers using IPv6 addresses.

### To create the pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
2. Click the **Create** button.
3. In the **Name** box, type a name for your pool. In our example, we type **UAG\_Pool\_IPv6**.
4. In the Health Monitors section, from the **Available** list, select **gateway-icmp**, and click the Add (<<) button.
5. From the Load Balancing Method list, select **Round Robin**. In this part of the configuration, a majority of the connections are not load balanced, they are sent to the appropriate UAG server based upon the entry in the connection table.
6. 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, type the IPv6 address of a UAG server. In our example, we type **2014:0:0:0:0:0:20**.
9. In the **Service Port** box, type **0**.
10. Click the **Add** button to add the member to the list.
11. Repeat steps 8-10 for each server you want to add to the pool. In our example, we repeat these steps once for the two remaining UAG servers, **2014:0:0:0:0:0:21** and **2014:0:0:0:0:0:22** (see Figure 19).
12. Click **Finished**.

Figure 19 Adding the IPv6 pool

## Creating the pre-selection iRule

The pre-selection iRule is designed to ensure that any traffic originating from any server to a remote client passes through the same UAG server to which the client has already attached. This ensures the tunnel between the client and the UAG server is reused for server originated traffic to the client.

### To create the 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. We use **linkpersist\_use**.
4. In the Definition section, copy and paste the following iRule:

```
rule testing {
  when CLIENT_ACCEPTED {
    set session_key [IP::local_addr]
    set forwardto [session lookup uie {$session_key any virtual}]

    if {$forwardto != ""}{
      nexthop internalnet $forwardto
      log local0. "I found the nexthop record for [IP::local_addr] and sent it to $forwardto"
    }
  }
}
```

- 
5. Click **Finished**.

## Creating the virtual server

After the pool and iRule have been created, we need to create the Virtual Server that takes new connections from the servers and sends them through the appropriate UAG servers.

### To create the virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this virtual server. In our example, we type **OUTBOUND\_VIP**.
4. In the Destination section, select the **Network** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we use **0.0.0.0**.
6. In the **Service Port** box, type **0**.
7. In the Configuration section, from the **Type** list, select **Performance (Layer 4)**.
8. From the **Protocol** list, select **UDP**.
9. In the Resources section, from the iRules **Available** list, select the name of the iRule you created in *Creating the pre-selection iRule*, on page 22, and click the Add (<<) button. In our example, we select **linkpersist\_use**.
10. From the **Default Pool** list, select the pool you created in *Creating the pool of UAG Servers (IPv6)*, on page 21. In our example, we select **UAG\_Pool\_IPv6**.
11. Click **Finished**.

This completes the deployment guide configuration.