DEPLOYMENT GUIDE
Version 1.4

IMPORTANT: This guide has been archived. While the content in this guide is still valid for the products and version listed in the document, it is no longer being updated and may refer to F5 or 3rd party products or versions that have reached end-of-life or end-of-support. See https://support.f5.com/csp/article/K11163 for more information.

Configuring IP Address Sharing in a Large Scale Network: DNS64/NAT64

Table of Contents

Configuring IP address sharing in a large scale network

Welcome to the F5 deployment guide for configuring IP address sharing in a large scale network. This document details how to configure the BIG-IP Local Traffic Manager (LTM) to address the looming IPv4 global address depletion dilemma, as well as for large organizations who want to deploy a large IPv4 private network or global IPv6 network internally.

While there are a number of possible ways to configure F5 devices to address these issues, this document focuses on two solutions that are straightforward and easy to implement. This guide is broken up into the following sections that give step-by-step guidance on how to configure the BIG-IP LTM for these solutions:

- *Configuring the BIG-IP LTM for the private IPv4 network*, on page 2
- *Configuring the BIG-IP LTM for global IPv6 with DNS64 and NAT64*, on page 7

For more information on the BIG-IP LTM, see ***www.f5.com/products/big-ip/product-modules/local-traffic-manager.html***.

Product versions and revision history

Product and versions tested for this deployment guide:

Product Tested	Version Tested
BIG-IP System	10.1

Revision history:

Document Version	Description
1.0	New deployment guide
1.1	Updated the full version of the iRule linked to on page 10 to properly rewrite the DNS compression pointer for the NS record's RDATA.
1.2	- Changed the virtual server Type to Performance (Layer 4) for the NAT64 virtual server on page 14. - Added guidance for Address translation and Port translation on the same virtual server. - Added a note about configuring an IPv4 Default Gateway for NAT64.
1.3	Removed the wildcard virtual server for outbound DNS traffic on page 13, as it was unnecesary.
1.4	Corrected the description of the iRule to choose a SNAT address (optional) on page 3.

Configuring the BIG-IP LTM for the private IPv4 network

It is well-known that private IP addresses (defined per RFC1918) have been used by most enterprises, some small service providers and mobile operators. The same private space could also be used to help overcome global IPv4 depletion.

In this scenario, the BIG-IP LTM acts as a gateway device that translates clients' private addresses to a set of global IPv4 addresses. This requires a wildcard forwarding virtual server with SNAT pool on the BIG-IP LTM.

Creating the SNAT Pool

In this procedure, we configure a SNAT pool. A secure network address translation (SNAT) translates the source IP address within a connection to a BIG-IP system IP address that you define. A SNAT pool is a group of these IP addresses.

Popular websites with heavy traffic (such as Google and Facebook) may require more SNAT addresses than typical websites. To estimate number of SNAT address you need for these high-traffic sites, you must have a number of SNAT addresses larger than the maximum number of concurrent connections per destination IP address divided by 64,000 (Number of SNAT address > (maximum concurrent connections per destination IP address / 64,000)).

For example, if the destination IP address has 250,000 maximum concurrent connections, you would need **4** SNAT addresses in the SNAT pool (250,000/64,000 = 3.906). You want to make sure you have enough SNAT addresses to cover the site with the highest expected traffic.

◆ Tip

It is outside the scope of this document to show you how to determine the number of concurrent connections, we recommend you use an appropriate monitoring tool. You could also use the BIG-IP LTM to determine concurrent connections by creating a virtual server that has a destination that matches the website you want to monitor. Then view the virtual server statistics to view the number of connections.
*Another indicator is if any of the log messages in **/var/log/ltm** mention port exhaustion.*
For more information, see the BIG-IP LTM documentation.

For more information on SNAT pools, see the *Configuring SNATs* chapter in the **Configuration Guide for Local Traffic Management**.

To create the SNAT pool

1. On the Main tab, expand **Local Traffic**, and then click **SNATs**.
2. On the Menu Bar, click **SNAT Pool List**.

3. Click the **Create** button.
4. In the **Name** box, type a name for this SNAT pool. In our example, we type **LSN-snat-pool**.
5. In the **IP Address** box, type an otherwise unused IP address, and click the **Add** button.

   Repeat this step for each additional address needed, from your calculation above.
6. Click **Finished**.

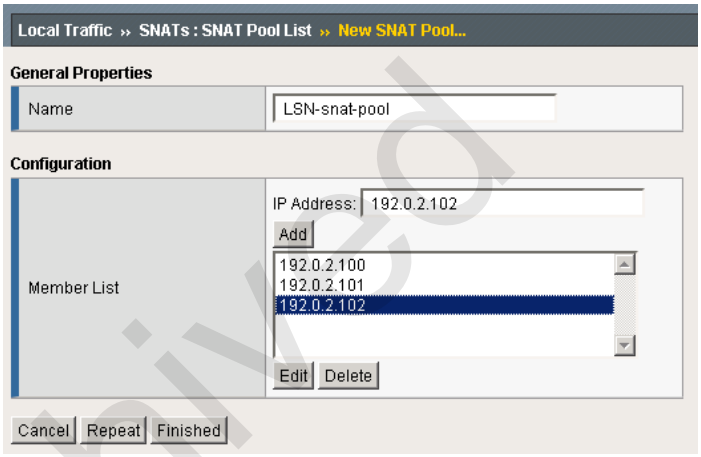

Figure 1 *SNAT pool configuration*

Using an iRule to choose a SNAT address (optional)

You can optionally define an iRule on the BIG-IP system that selects a specific SNAT address. By default, the BIG-IP LTM picks address from the SNAT pool in a round-robin fashion. Some applications or websites may use multiple connections (or even multiple destinations) per session, and the application may break if the client accesses from a different source IP. The following simple iRule can be used to ensure client always uses the same SNAT address.

This iRule picks a SNAT address related to the last octet of the client IP address. For example, if the client address is 172.16.33.57, the BIG-IP system uses the SNAT address 10.10.10.58 (notice the last octet is one higher). This simple example assumes there are 64 SNAT addresses from 10.10.10.1 through 10.10.10.64.

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 type **snat-pool-irule**.

4. In the Definition section, copy and paste the following iRule:

```
1   when CLIENT_ACCEPTED {
2       snat 10.10.10.[expr ( [getfield [IP::client_addr] "." 4] % 64 ) + 1 ]
3   }
```

5. Click **Finished**.

Note

This SNAT iRule overwrites the SNAT setting from the virtual server.

Creating the wildcard virtual server

The next task it to configure a wildcard virtual server that contains the SNAT pool you created.

By default, this wildcard virtual server is enabled on all VLANs. As an optional step for added security, you can lock down the virtual server to specific VLANs in step 11.

To create the wildcard 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. We type **LSN_wildcard**.
4. In the Destination row, click 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 ***** or select ***All Ports** from the list.
8. From the **Configuration** list, select **Advanced**.
9. From the **Type** list, select **Forwarding (IP)**.
10. From the **Protocol** list, select ***All Protocols**.
11. *Optional*: From the **VLAN and Tunnel Traffic** (or **VLAN Traffic** in some versions) list, select **Enabled on**.
    From the Available list, select the appropriate VLANs and then click the Add (<<) button.
12. From the **SNAT Pool** list, select the SNAT Pool you created in *Creating the SNAT Pool*. In our example, we select **lns-snat-pool**.

13. *Optional*: If you created the iRule in *Using an iRule to choose a SNAT address (optional)*, in the Resources section, from the **iRule Available** list, select the iRule you created and click the Add (<<) button.
14. Click **Finished**.

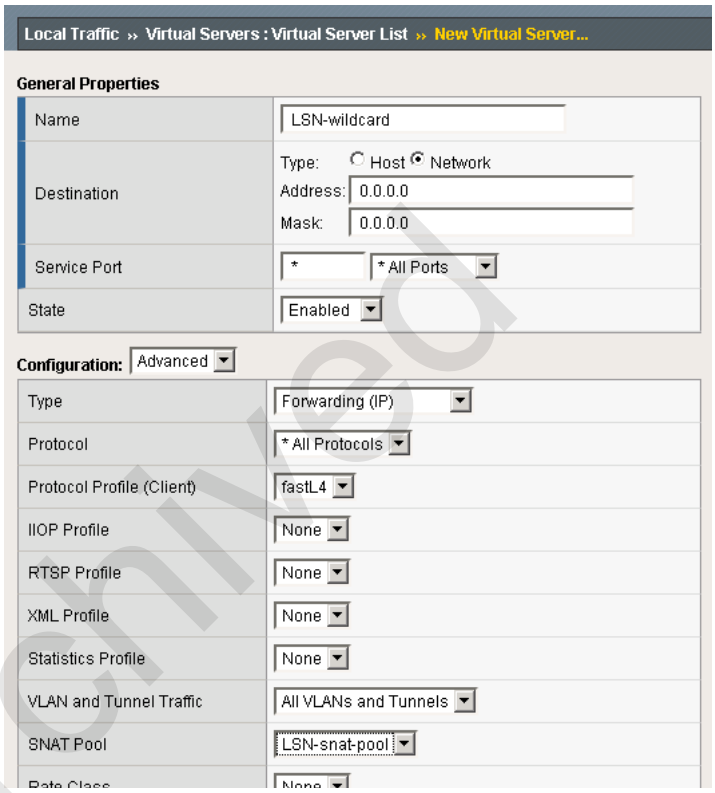

Figure 2 *Wildcard virtual server configuration (truncated)*

Note

Because this is a forwarding virtual server, the BIG-IP system forwards traffic based on routing. You may need to configure a default gateway on the BIG-IP LTM. See the online help or BIG-IP documentation for more information.

Supporting NAT with active FTP mode (optional)

To support active mode FTP traffic from clients, you need to configure an additional wildcard virtual server and apply an FTP profile to it. We create this virtual server so the BIG-IP system can detect the ephemeral data port which the client opens, and rewrite it to the SNAT address/available

ephemeral port. The BIG-IP system also listens on the address/port (that it rewrites) and translates it to the actual address/port to which the client is listening.

The first task in this section is to create a FTP profile.

To create the FTP 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 **Services** menu, click **FTP**.
3. Click the **Create** button.
4. In the **Name** box, type a name. In our example, we type **LSN-ftp**.
5. Configure any of the settings as applicable for your configuration. In our example, we leave the defaults.
6. Click **Finished**.

Next we create the virtual server.

To create the wildcard 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. We type **FTP_wildcard**.
4. In the Destination row, click 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 **21** or select **FTP** from the list.
8. Leave the **Type** list set to **Standard**.
9. From the **FTP Profile** list, select profile you created in the preceding procedure. In our example, we select **LSN-ftp**.
10. *Optional*: From the **VLAN and Tunnel Traffic** (or **VLAN Traffic** in some versions) list, select **Enabled on**.
    From the Available list, select the appropriate VLANs and then click the Add (<<) button.
11. From the **SNAT Pool** list, select the SNAT Pool you created in *Creating the SNAT Pool*. In our example, we select **LNS-snat-pool**.
12. Click **Finished**.

This concludes the BIG-IP LTM configuration for the private IPv4 network.

Configuring the BIG-IP LTM for global IPv6 with DNS64 and NAT64

In this section, we configure the BIG-IP LTM for global IPv6 with DNS64 and NAT64.

With this solution, clients are in IPv6 network, and the BIG-IP system should be on both the IPv6 and IPv4 networks (see Figure 3).

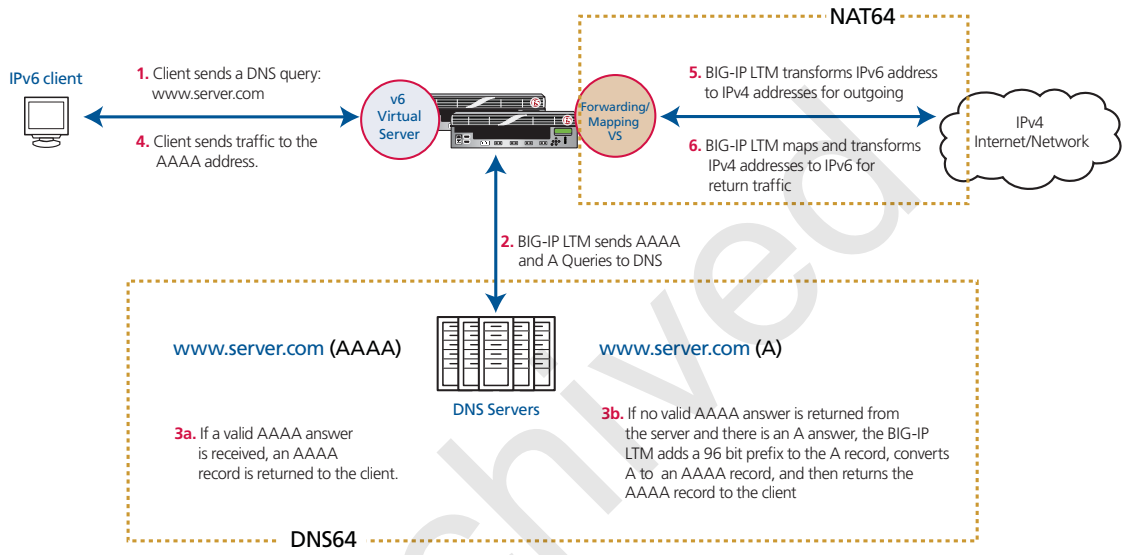

Figure 3 *Logical configuration example and flow*

For this solution, the BIG-IP system should be on both IPv6 and IPv4 network as shown.

Note

The IPv6 network may or may not connect to the global IPv6 network. Also, IPv6 clients maybe able to access full or parts of the global IPv6 network.

When a client initiates access the Internet, the process typically begins by sending a DNS request for the IP address of the destination they are trying to access. The DNS requests can take one of two forms:

- **IPv6**
  The DNS query from an IPv6 client is called an **AAAA** query. The client uses the AAAA query to request an IPv6 address that is associated with the name.
- **IPv4**
  The DNS query from an IPv4 client is called an **A** query. The client uses the A query to request an IPv4 address that is associated with the name.

See RFC3596 for more information: ***http://www.ietf.org/rfc/rfc3596.txt***.

Important

You must configure the client or the network so that the DNS requests are routed to BIG-IP. Consult the appropriate documentation.

Once the BIG-IP receives an AAAA DNS query, it forwards it to the DNS server. If the DNS server returns the AAAA answer with the associated IPv6 address record, the BIG-IP system forwards the answer back to the IPv6 client.

If the server responds that there is no IPv6 address associated to this name, BIG-IP issues another DNS query, but this time an A query.

If the DNS server returns an A answer that there is no IPv4 address associated with this name, the BIG-IP transforms the A answer to an AAAA answer and forwards it to the IPv6 client.

If the DNS server returns an A answer with an associated IPv4 address record, the BIG-IP creates the IPv6 address record from the IPv4 address record, transforms the A answer to an AAAA answer, including the synthesized IPv6 address, and sends it back to IPv6 client.

The easiest way to create an IPv6 address from an IPv4 address is to put a 96-bit prefix in front of the 32-bit IPv4 address (IPv6 addresses contain 128 bits). The 96-bit prefix could be an unused routable /96 IPv6 subnet. ***Important:*** The traffic destined to this /96 IPv6 subnet must be routed to BIG-IP system.

You can also find recommendations regarding /96 bit prefixes from the following IETF dns64 draft document: ***https://datatracker.ietf.org/doc/draft-ietf-behave-dns64/***.

This section is divided into two parts: *Configuring the BIG-IP LTM for DNS64*, and *Configuring the BIG-IP LTM for NAT64*, on page 12.

Configuring the BIG-IP LTM for DNS64

In this section, we configure the BIG-IP LTM for DNS64.

Important

This configuration assumes the DNS servers have connectivity to the Internet.

Creating a DNS health monitor

In the following procedure, we use the DNS Application Template to create the DNS health monitor. The DNS Application Template is only available in BIG-IP LTM versions 10.2 and later.

To run the generic DNS application template

1. Verify that your current administrative partition is set to **Common**. The Partition list is in the upper right corner.

2. On the Main tab, expand **Templates and Wizards**, and then click **Templates**. The Templates screen opens.
3. In the Application column, click **Generic DNS**.
4. In the **Virtual Server Questions** section, complete the following:
   a) You can type a unique prefix for the DNS objects this template creates. In our example, we leave this setting at the default, **my_DNS_**.
   b) Select **Yes** from the list asking if you only want the template to configure the monitor.
5. In the **Monitor Questions** section, complete the following:
   a) From the Record type list, select the record type you want to use to test the DNS servers. In our example, we leave this at the default **A**.
   b) Type the host name you want to send the DNS server for this monitor. In our example, we type **www.example.com**.
   c) Type IP address you expect to receive from the DNS server from the host name in step b. In our example, we type **192.0.2.50**.
6. Click the **Finished** button.

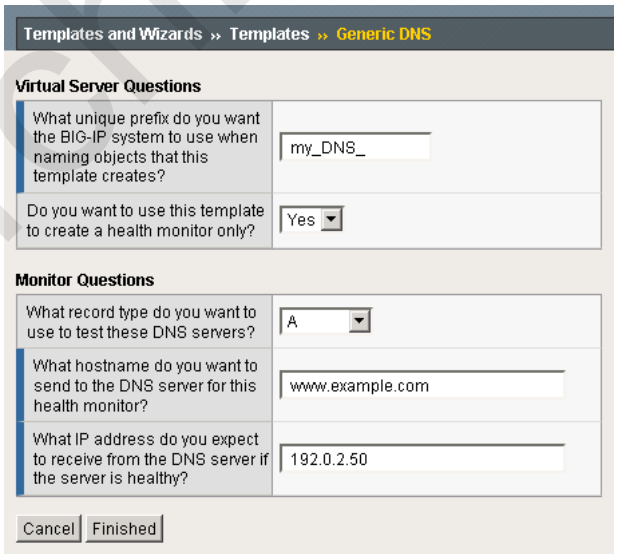

Figure 4 *DNS monitor configuration*

Creating the DNS pool

Next, we create the load balancing pool that contains the DNS servers. The servers can have IPv4 or IPv6 addresses.

To create the pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
2. Click the **Create** button. The New Pool screen opens.
3. In the **Name** box, type a name for your pool. We type **DNS-pool**.
4. In the **Health Monitors** section, select the name of the monitor that was created by the template in *Creating a DNS health monitor* (or the transparent monitor if you created one), and click the Add (**<<**) button. In our example, we select **my_DNS__monitor**.
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 **Least Connections (member)**.
6. In this pool, we leave the Priority Group Activation **Disabled**.
7. In the New Members section, select the **New Address** option.
8. In the **Address** box, add the first DNS server to the pool. In our example, we type **10.132.81.100**.
9. In the **Service Port** box, type **53**.
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 four times for the remaining servers, **10.132.81.101 - .104**.
12. Click the **Finished** button.

Creating the iRule

The next task is to create an iRule for the DNS traffic. There are two options for this iRule, choose the version appropriate for your configuration. For convenience, the iRule code is available in downloadable text files.

- *Simple version*
  This version of the iRule assumes the IPv6 client does not have a route to any global IPv6 network (IPv6 Internet). In this case, the AAAA query is transformed to an A query.
  Download:
  http://www.f5.com/solutions/resources/deployment-guides/files/dns64_simple.txt
- *Full version*
  This is the complete version of the iRule that sends out an AAAA query first before sending an A query.
  Download:
  http://www.f5.com/solutions/resources/deployment-guides/files/dns64_full.txt

Be sure to download the appropriate iRule before performing the following procedure.

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 type **DNS64-irule**.
4. In the Definition section, copy and paste the iRule you downloaded:
5. ***Important***: Locate the following lines in the iRule, and change the 96-bit prefix to a prefix applicable to your configuration:

```
when RULE_INIT {
    set static::prefix "20020123000000000000000"
}
```

6. Click **Finished**.

Creating a SNAT pool

The next task is to create a SNAT Pool. To create the SNAT pool, use the procedure *Creating the SNAT Pool*, on page 2.

Creating the virtual server

The next task is to create the virtual server. For this virtual server, configuring a SNAT pool is optional. If your configuration requires a SNAT pool, before configuring this virtual server, see *Creating the SNAT Pool*, on page 2 before configuring this 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 **dns-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 **2001:0db8:85a3:08d3:1319:8a2e:0370:7334**.
6. In the **Service Port** box, type **53**.
7. From the **Configuration** list, select **Advanced**.
8. From the **Protocol** list, select **UDP** (see Figure 5).

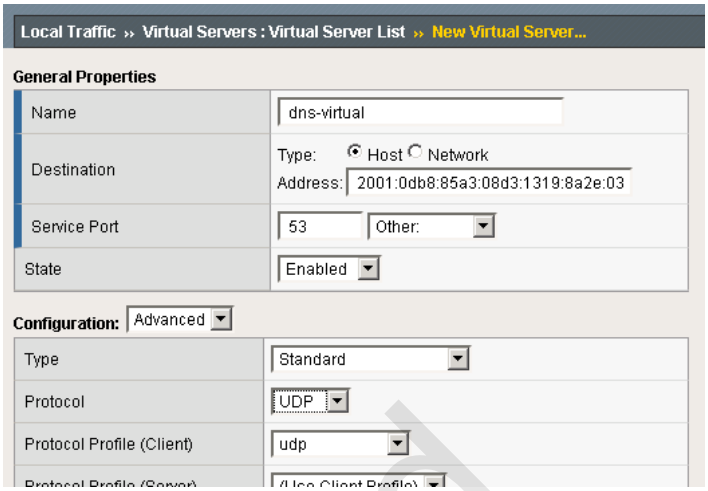

Figure 5 *IPv6 virtual server configuration (truncated)*

9. *Optional:* From the **SNAT Pool** list, select the SNAT Pool you created in *Creating a SNAT pool*. In our example, we select **lns-SNAT-pool**.
10. *Optional*: From the **VLAN and Tunnel Traffic** (or **VLAN Traffic** in some versions) list, select **Enabled on**.
From the Available list, select the appropriate VLANs and then click the Add (<<) button.
11. In the Resources section, from the **iRule Available** list, select the iRule you created in *Creating the iRule*, on page 10 and click the Add (<<) button.
12. From the **Default Pool** list, select the pool you created in *Creating the DNS pool*, on page 9.
13. Click the **Finished** button.

Configuring the BIG-IP LTM for NAT64

In this section, we configure the BIG-IP system for NAT64. NAT64 is a common solution put in place to overcome IPv4 depletion. NAT64 is simply a NAT, but it is actually NAT'ing all IPv4 addresses to an IPv6 subnet (with a /96 mask).

◆ Note

Your BIG-IP LTM must have an IPv4 Default Gateway that allows the LTM to reach any address behind it. In our example, we would have an IPv4 Default Gateway of 192.0.2.254. For information on configuring a Default Gateway on the BIG-IP LTM, see the BIG-IP LTM documentation. You can also see the online help (on the Main tab, expand Network, click Routes, click the Add button, and then click the Help tab).

Creating a SNAT pool

The first task is to create a SNAT Pool. To create the SNAT pool, use the procedure *Creating the SNAT Pool*, on page 2.

Creating the translation iRule

The BIG-IP system acts as a gateway between the IPv6 and IPv4 networks. When IPv6 clients send out a packet, the BIG-IP translates both source IP and destination IP; the source IP translation is handled by the SNAT pool, and the destination IP address translation is handled by the following iRule.

This iRule takes the last 32 bits from IPv6 destination address and uses that as a new IPv4 destination address when it forwards this packet out on IPv4 network.

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 type **NAT64-irule**.
4. In the Definition section, copy and paste the following iRule:

```
1   when CLIENT_ACCEPTED {
2       node [string range [IP::addr [IP::local_addr] mask ::ffff:ffff] 0 end]
3   }
```

5. Click **Finished**.

Optional: You can optionally add a line to the iRule that selects a specific SNAT address. By default, the BIG-IP LTM picks address from the SNAT pool in a round-robin fashion. Some applications or websites may use multiple connections (or even multiple destinations) per session, and the application may break if the client accesses from a different source IP. The following simple iRule can be used to ensure client always uses the same SNAT address. This iRule picks a SNAT address by using the last octet of the client IP address. For example, if client address is 172.16.33.57, the BIG-IP system uses SNAT address 10.10.10.57. This simple example assumes there are 64 SNAT addresses starting from 10.10.10.1 to 10.10.10.64.

To use this optional line, step 4 in the preceding procedure would be:

4. In the Definition section, copy and paste the following iRule:

```
1   when CLIENT_ACCEPTED {
2     node [string range [IP::addr [IP::local_addr] mask ::ffff:ffff] 0 end]
3     snat 10.10.10.[expr ( [getfield [IP::addr [IP::client_addr] mask ::ff] "." 4] % 64 ) + 1 ]
4   }
```

Creating the IPv6 network virtual server

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

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. We type **NAT64-virtual**.
4. In the Destination section, click the **Network** option button.
5. In the **Address** box, type the appropriate IPv6 address based on the 96-bit prefix. For example, if 96-bit prefix is 2002:0123:0000:0000:0000:0000, the IP address should be 2002:0123:0000:0000:0000:0000::
6. In the **Mask** box, type the associated subnet mask. Using the Address in the example above, we type **ff:ff:ff:ff:ff:ff::** for the mask.
7. In the **Service Port** box, type ***** or select ***All Ports** from the list.
8. From the **Type** list, select **Performance (Layer 4)**.
9. From the Protocol list, select ***All Protocols**.
10. From the **Address Translation** row, make sure the **Enabled** box is checked to enable Address Translation (the default).
11. From the **Port Translation** row, clear the check from the **Enabled** box to ***disable*** Port Translation.
12. In the Resources section, from the **iRule Available** list, select the iRule you created in *Creating the translation iRule*, on page 13. In our example, we select **NAT64-irule**.
13. Click **Finished**.

Supporting NAT with active FTP mode (optional)

You can optionally configure the BIG-IP LTM to support NAT with active FTP mode.

To configure support for active FTP, follow *Supporting NAT with active FTP mode (optional)*, on page 5 and 6 with the following exceptions:

- *Step 5 of creating a virtual server*: In the Destination section, click the **Network** button. In the **Address** box, type the appropriate IPv6 destination. In our example, we type **2002:123::**.
- *Step 6 of creating a virtual server*: In the **Mask** box, type the appropriate IPv6 mask. In our example, we type **ff:ff:ff:ff:ff:ff::**

- *Additional step*: In the Resources section, from the **iRule Available** list, select the iRule you created in *Creating the translation iRule*, on page 13. In our example, we select **NAT64-irule**.

This concludes the configuration.