Deploying the BIG-IP System for nPath routing

Welcome to the F5® deployment guide for nPath routing. This document contains guidance on configuring the BIG-IP® system version 11.4 and later for nPath routing (also known as Asymmetric routing or Direct Server Return (DSR)). BIG-IP version 11.0 introduced iApp™ Application templates, an extremely easy way to accurately configure the BIG-IP system for your nPath deployment.

Products and Versions tested

Product	Version
BIG-IP LTM	11.4, 11.4.1, 11.5, 11.5.1, 11.6
nPath	Not applicable
nPath iApp template	System iApp that ships with v11.4 and later
Deployment Guide version	1.2 (see *Document Revision History on page 18)*

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

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

Contents

What is nPath routing?

With the nPath routing configuration, you can route outgoing server traffic around the BIG-IP system directly to an outbound router. This method of traffic management increases outbound throughput because packets do not need to be transmitted to the BIG-IP system for translation and forwarding to the next hop.

In bypassing the BIG-IP system on the return path, nPath routing departs significantly from a typical load-balancing configuration. In a typical load-balancing configuration, the destination address of the incoming packet is translated from that of the virtual server to that of the server being load balanced to, which then becomes the source address of the returning packet. A default route set to the BIG-IP system then sees to it that packets returning to the originating client return through the BIG-IP system, which translates the source address back to that of the virtual server.

For more information on nPath routing, see the BIG-IP documentation.

What is F5 iApp™?

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

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

Prerequisites and configuration notes

The following are general prerequisites and configuration notes for this guide:

- For this deployment guide, the BIG-IP system ***must*** be running version 11.4 or later. If you are using a previous version of the BIG-IP system, see the deployment guide index on F5.com. The configuration in this guide does not apply to previous versions.
- If you upgraded your BIG-IP system from a previous v11 version, and have an existing Application Service that used the f5.npath iApp template, see *Upgrading an Application Service from previous version of the iApp template on page 12*.
- This document provides guidance for using the iApp for nPath found in version 11.4 and later. For users familiar with the BIG-IP system, there is a manual configuration table at the end of this guide. However, we recommend using the iApp template.

Using this guide

This guide is intended to help users deploy web-based applications using the BIG-IP system. This deployment guide contains guidance on two ways to configure the BIG-IP system: using the iApp template, and manually configuring the BIG-IP system.

Using this guide to configure the App template

We recommend using the iApp template to configure the BIG-IP system for your nPath routing implementation. The majority of this guide describes the iApp template and the different options the template provides for configuring the system for nPath routing.

The iApp template configuration portion of this guide walks you through the entire iApp, giving detailed information not found in the iApp or inline help. The questions in the iApp template itself are all in a table and at the same level. In this guide, we have grouped related questions and answers in a series of lists. Questions are part of an ordered list and are underlined and in italics or bold italics. Options or answers are part of a bulleted list, and in bold. Questions with dependencies on other questions are shown nested under the top level question, as shown in the following example:

1. ***Top-level question found in the iApp template***
   - *Select an object you already created from the list* (such as a profile or pool; not present on all questions. Shown in bold italic)
   - **Choice #1** (in a drop-down list)
   - **Choice #2** (in the list)
     a. *Second level question dependent on selecting choice #2*
        - **Sub choice #1**
        - **Sub choice #2**
          i). *Third level question dependent on sub choice #2*
              - **Sub-sub choice**
              - **Sub-sub #2**
                1). *Fourth level question (rare)*

Manually configuring the BIG-IP system

Users already familiar with the BIG-IP system can use the manual configuration tables to configure the BIG-IP system for the nPath implementation. These configuration tables only show the configuration objects and any non-default settings recommended by F5, and do not contain procedures on specifically how to configure those options in the Configuration utility. See *Appendix: Manual configuration table on page 14.*

Preparing to use the iApp

In order to use the iApp for nPath routing, it is helpful to have some information, such as server IP addresses and domain information before you begin. Use the following table for information you may need to complete the template. The table does not contain every question in the template, but rather includes the information that is helpful to have in advance. More information on specific template questions can be found on the individual pages.

<table>
<tr><th colspan="3">BIG-IP LTM Preparation table</th></tr>
<tr><td>Basic/Advanced mode</td><td colspan="2">In the iApp, you can configure your implementation with F5 recommended settings (Basic mode) which are a result of testing the nPath configuration. Advanced mode gives you the to configure the BIG-IP system on a much more granular level, configuring specific options, or using your own pre-built profiles or iRules. Basic and Advanced "configuration mode" is independent from the Basic/Advanced list at the very top of the template which only toggles the Device and Traffic Group options (see page 6)</td></tr>
<tr><td rowspan="3">Virtual Server and Pools</td><td>Virtual Server</td><td>Pool</td></tr>
<tr><td>The Virtual server is the address clients use to access the servers.</td><td>The load balancing pool is the LTM object that contains the servers.</td></tr>
<tr><td>IP address for the virtual server:

Associated service port:</td><td>IP addresses of the servers: 1:
2: 3:
4: 5:
6: 7:
8: 9:</td></tr>
<tr><td rowspan="2">Profiles</td><td colspan="2">For each of the following profiles, the iApp will create a profile using the F5 recommended settings (or you can choose ‘do not use’ many of these profiles). While we recommend using the profiles created by the iApp, you have the option of creating your own custom profile outside the iApp and selecting it from the list. The iApp gives the option of selecting our the following profiles (some only in Advanced mode). Any profiles must be present on the system before you can select them in the iApp</td></tr>
<tr><td colspan="2">Persistence (optional) | FastL4</td></tr>
<tr><td>iRules</td><td colspan="2">In Advanced mode, you have the option of attaching iRules you create to the virtual server created by the iApp. For more information on iRules, see https://devcentral.f5.com/irules Any iRules you want to attach must be present on the system at the time you are running the iApp.</td></tr>
</table>

Configuring the BIG-IP iApp for nPath routing

Use the following guidance to help configure the BIG-IP system for nPath routing using the BIG-IP iApp template.

Getting Started with the iApp for nPath routing

To begin the nPath routing iApp Template, use the following procedure.

1. Log on to the BIG-IP system.
2. On the Main tab, expand **iApp**, and then click **Application Services**.
3. Click **Create**. The Template Selection page opens.
4. In the **Name** box, type a name. In our example, we use **nPath-iapp_.**
5. From the **Template** list, select **f5.npath**. The nPath template opens.

Advanced options

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

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

Template Options

This section contains general questions about the way you configure the iApp template.

1. ***Do you want to see inline help?***
   Choose whether you want to see informational and help messages inline throughout the template, or if you would rather hide this inline help. If you are unsure, we recommend having the iApp display the inline help.
   Important and critical notes are always shown, no matter which selection you make.
   - **Yes, show inline help text**
     Select this option to see all available inline help text.
   - **No, do not show inline help text**
     If you are familiar with this iApp template, or with the BIG-IP system in general, select this option to hide the inline help text.

2. ***Which configuration mode do you want to use?***
   Select whether you want to use F5 recommended settings, or have more granular, advanced options presented.
   - **Basic - Use F5's recommended settings**
     In basic configuration mode, options like load balancing method and parent profiles are all set automatically. The F5 recommended settings come as a result of extensive testing with web applications, so if you are unsure, choose Basic.
   - **Advanced - Configure advanced options**
     In advanced configuration mode, you have more control over individual settings and objects, such as server-side optimizations and advanced options like Slow Ramp Time and Priority Group Activation. You can also choose to attach iRules you have previously created to the application service. The Advanced option provides more flexibility for experienced users.

     Advanced options in the template are marked with the Advanced icon: Advanced. If you are using Basic/F5 recommended settings, you can skip the questions with this icon.

Network

This section contains questions about your networking configuration. This whole section only appears if you selected Advanced mode.

1. ***Which VLANs transport client traffic?*** Advanced
   The BIG-IP system allows you to restrict client traffic to specific BIG-IP VLANs, which can provide an additional layer of security, as only traffic from the VLANs you select are allowed to the servers. By default, all VLANs configured on the system are enabled. Select which of your BIG-IP VLANs are transporting client traffic. If you want the BIG-IP system to only accept client traffic from specific VLANs, from the **Selected** list, select the appropriate VLAN(s) from which you do not want the system to accept traffic, and then click the Remove (**>>**) button to move the VLAN to the Option box.

   Note

   If you choose to allow traffic from certain VLANs, when additional VLANs are added to the BIG-IP system at a later time, this iApp configuration will deny traffic from these VLANs by default. To accept traffic from these VLANs, you must re-enter the template and add the VLAN(s).

High Availability

This section gathers information about your nPath deployment that will be used in the BIG-IP *virtual server* and *load balancing pool*.

1. ***What IP address do you want to use for the virtual server?***
   Type the IP address you want to use for the BIG-IP virtual server. This is the address clients use (or a DNS entry resolves to this address) to access the deployment via the BIG-IP system.

2. ***What is the associated service port?***
   Type the port number you want to use for the BIG-IP virtual server. This depends on the application you are using.

3. ***What protocol do you want this virtual server to process?***
   Choose the network protocol you want the system to use to direct traffic on this virtual server. If you are unsure, we recommend choosing ***All Protocols**.

   - **TCP**
     Selecting this option specifies that the virtual server supports the TCP protocol.
   - **UDP**
     Selecting this option specifies that the virtual server supports the UDP protocol.
   - ***All protocols**
     Selecting this option specifies that the virtual server supports all network protocols.

4. ***How long a timeout (in seconds) do you want before closing unused connections?***
   Type the number of seconds you want to use before the BIG-IP system closes unused connections. We recommend the default of 51 seconds.

5. ***Which persistence profile do you want to use?*** Advanced
   Choose whether you want connections to persist to the same server after the initial connection has been made.

   Creating a custom profile is not a part of this template; see **Local Traffic** >> **Profiles** : **Persistence** to create a persistence profile. To select any new profiles you create, you need to restart or reconfigure this template.

   Select one of the following persistence options:

   - **Do not use persistence**
     Select this option if you do not require persistence
   - *Select an existing persistence profile*
     If you have already created a persistence profile for this configuration, you can select it from the list.

6. ***What Fast L4 profile do you want to use?*** Advanced
The purpose of a Fast L4 profile is to help you manage Layer 4 traffic more efficiently. When you assign a Fast L4 profile to a virtual server, the Packet Velocity ASIC (PVA) hardware acceleration within the BIG-IP system (if supported) can process some or all of the Layer 4 traffic passing through the system. By offloading Layer 4 processing to the PVA hardware acceleration, the BIG-IP system can increase performance and throughput for basic routing functions (Layer 4) and application switching (Layer 7)

Unless you have requirements for configuring specific settings, we recommend allowing the iApp to create a new profile. Creating a custom profile is not a part of this template; see **Local Traffic** >> **Profiles** : **Protocol : Fast L4** to create a Fast L4 profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Use F5's recommended Fast L4 profile**
Chose this option to enable the iApp to create the F5 recommended Fast L4 profile.
- *Select an existing Fast L4 profile*
If you have previously created a Fast L4 profile, you have the option of selecting it instead of allowing the iApp to create a new one. From the list, select an existing profile.

7. ***Do you want to create a new pool or use an existing one?***
A *load balancing pool* is a logical set of servers, grouped together to receive and process traffic. When clients attempt to access the servers via the BIG-IP virtual server, the BIG-IP system distributes requests to any of the servers that are members of that pool.

- *Select an existing pool*
If you have already created a pool for your servers, you can select it from the list.
If you do select an existing pool, all of the rest of the questions in this section disappear.
- **Create a new pool**
Leave this default option to create a new load balancing pool and configure specific options.

a. *Which load balancing method do you want to use?* Advanced
Specify the load balancing method you want to use for this pool. We recommend the default, **Least Connections (member)**.

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

> **Important**
>
> *TCP Request Queuing is an advanced feature and should be used only if you understand how it will affect your deployment, including application behavior and BIG-IP performance.*
> *If you enable TCP Request Queuing, you must have a Connection Limit set on at least one of the nodes when configuring the Address/Port.*

- **Do not enable TCP request queuing**
Select this option if you do not want the BIG-IP system to queue TCP requests.
- **Enable TCP request queuing**
Select this option if you want to enable TCP request queuing on the BIG-IP system.

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

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

c. *Use a Slow Ramp time for newly added servers?* **Advanced**
With Slow Ramp, the BIG-IP system gradually adds connections to a newly-enabled or newly-added server over a time period you specify, rather than sending a full proportion of the traffic immediately. Slow Ramp is essential when using load balancing methods like Least Connections, as the system would otherwise send all new connections to a new server immediately, potentially overwhelming that server.

Select whether you want to use a Slow Ramp time.

- **Use Slow Ramp**
  Select this option for the system to implement Slow Ramp time for this pool.
  i). *How many seconds should Slow Ramp time last?*
  Specify a duration in seconds for Slow Ramp. The time period you select for Slow Ramp is highly dependent on the speed of your server hardware and the behavior of your services. The default setting of 300 seconds (5 minutes) is very conservative in most cases.
- **Do not use Slow Ramp**
  Select this option if you do not want to use Slow Ramp. If you select this option, we recommend you do not use the Least Connections load balancing method.

d. *Do you want give priority to specific groups of servers?* **Advanced**
Select whether you want to use Priority Group Activation. Priority Group Activation allows you to segment your servers into priority groups. With Priority Group Activation, the BIG-IP system load balances traffic according to the priority number you assign to the pool members. A higher number indicates higher priority. Traffic is only sent to the servers with the highest priority, unless the number of available servers in that priority group falls below the value you specify as the minimum. The BIG-IP system then sends traffic to the group of servers with the next highest priority, and so on. See the BIG-IP documentation for more details.

- **Do not use Priority Group Activation**
  Select this option if you do not want to enable Priority Group Activation.
- **Use Priority Group Activation**
  Select this option if you want to enable Priority Group Activation.
  You must add a priority to each server in the Priority box described in #e.
  i). *What is the minimum number of active members in a group?*
  Specify the minimum number of servers that must be active to continue sending traffic to the priority group. If the number of active servers falls below this minimum, traffic will be sent to the group of servers with the next highest priority group number.

e. *Which servers are a part of this pool?*
Specify the IP address(es) of your servers. If you have existing nodes on this BIG-IP system, you can select them from the list, otherwise type the addresses. You can optionally add a Connection Limit. If you enabled Priority Group Activation, you must also specify a Priority for each device. Click **Add** to include additional servers.

Application Health

In this section, you answer questions about how you want to implement application health monitoring on the BIG-IP system.

1. ***Create a new health monitor or use an existing one?***
Application health monitors are used to verify the servers are available and functioning.

Unless you have requirements for configuring other options not in the following list of questions, we recommend allowing the iApp to create a new monitor. Creating a custom health monitor is not a part of this template; see **Local Traffic** >> **Monitors**. To select any new monitors you create, you need to restart or reconfigure this template.

- *Select the monitor you created from the list*
  If you manually created the health monitor for your servers, select it from the list. Continue with the next section.

- **Create a new ICMP monitor**
  Select this option if you want the template to create a simple ICMP ping monitor. The server is marked up if the ping is successful.

  a. *How many seconds between health checks?*
  Specify how long the system should wait between each health check. This is used as the Interval setting for the monitor. We recommend the default of 30 seconds.

- **Create a new TCP monitor**
  Select this option if you want the iApp to create a new TCP monitor, and then answer the following questions:

  a. *How many seconds between health checks?*
  Specify how long the system should wait between each health check. This is used as the Interval setting for the monitor. We recommend the default of 30 seconds.

  b. *What string should be sent to verify server health?*
  You can configure the template to retrieve a specific page by typing the path here. This allows a much more granular health check. Leaving the default marks the node up if anything is returned from the web page.

  c. *What string should the system expect as a response?*
  If you entered a unique request in the previous question, this is where you enter the expected response from that query. The node is marked up only if this specific response is received.

- **Create a new UDP monitor**
  Select this option if you want the iApp to create a new UDP monitor, and then answer the following questions:

  a. *How many seconds between health checks?*
  Specify how long the system should wait between each health check. This is used as the Interval setting for the monitor. We recommend the default of 30 seconds.

  b. *What string should be sent to verify server health?*
  You can configure the template to retrieve a specific page by typing the path here. This allows a much more granular health check. Leaving the default marks the node up if anything is returned from the web page.

  c. *What string should the system expect as a response?*
  If you entered a unique request in the previous question, this is where you enter the expected response from that query. The node is marked up only if this specific response is received.

iRules

In this section, you can add custom iRules to the nPath deployment. This entire section is available only if you selected Advanced mode. Because the iApp template creates two virtual servers, one for TCP traffic and one for UDP traffic, there are separate questions for attaching iRules to the individual virtual server.

iRules are a scripting language that allows an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. An iRule contains the set of instructions the system uses to process data flowing through it, either in the header or payload of a packet.

Warning

While iRules can provide additional functionality not present in the iApp, iRules are an advanced feature and should be used only if you understand how each iRule will affect your deployment, including application behavior and BIG-IP system performance.

1. ***Do you want to add any custom iRules to the virtual server?*** Advanced
   Select if have preexisting iRules you want to add to the nPath implementation.

   If you have iRules you want to attach to the TCP virtual server the iApp creates for your servers, from the **Options** box, click the name of the applicable iRule(s) and then click the Add (**<<**) button to move them to the **Selected** box.

Finished

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

Next steps

After completing the iApp Template, the BIG-IP Application Services page opens for the nPath service you just created. To see the list of all the configuration objects created to support nPath routing, on the Menu bar, click **Components**. The complete list of all related objects opens. You can click individual objects to see the settings.

Once the objects have been created, you are ready to use the new deployment.

Modifying DNS settings to use the BIG-IP virtual server address

Before sending traffic to the BIG-IP system, your DNS administrator may need to modify any DNS entries for the implementation to point to the BIG-IP system's virtual server address.

Modifying the iApp configuration

The iApp application service you just created can be quickly and easily modified if you find it necessary to make changes to the configuration. The Strict Updates feature of the iApp prevents users from manually modifying the iApp configuration (Strict Updates can be turned off, but use extreme caution). iApp allows you to re-enter the template, make changes, and then update the template. The modifications are automatically made to any of the associated objects.

To modify the configuration

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. Click the name of your nPath Application Service from the list.
3. On the Menu bar, click **Reconfigure**.
4. Make the necessary modifications to the template.
5. Click the **Finished** button.

Viewing statistics

You can easily view a number of different statistics on the BIG-IP system related to the configuration objects created by the iApp template. You can get statistics specific to the Application Service if you have provisioned AVR. Otherwise, you can always get object-level statistics.

Object-level statistics

Use the following procedure to view statistics.

To view object-level statics

1. On the Main tab, expand **Overview**, and then click **Statistics**.
2. From the **Statistics Type** menu, you can select **Virtual Servers** to see statistics related to the virtual servers.
3. You can also choose **Pools** or **Nodes** to get a closer look at the traffic.
4. To see Networking statistics in a graphical format, click **Dashboard**.

For more information on viewing statistics on the BIG-IP system, see the online help or product documentation.

Upgrading an Application Service from previous version of the iApp template

If you upgraded your BIG-IP system from a previous v11 version and had an existing Application Service that used the f5.npath template from one of those versions, you will see a warning that the source template has changed. In version 11.4 and later, the f5.npath template has been significantly improved, and we strongly recommend you upgrade the source template to the new template available in v11.4 and later.

When you upgrade to the current template version, the iApp retains all of your settings for use in the new template. You will notice the location of the questions are different in the new version of the template, most questions are asked in a different way, and BIG-IP WebAccelerator is now called BIG-IP Application Acceleration Manager. There are also many more options you can configure in the new version of the template.

To upgrade an Application Service to the current version of the template

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. From the list, click the name of the application service you created using the f5.npath template. You'll see a warning icon in the Template Validity column.
3. On the Menu bar, click **Reconfigure**.
4. In the Template Options section, from the **Do you want to upgrade this template** question, select **Yes**.
5. Without changing any settings, click the **Finished** button. The system creates an application service object with only the new template object in the Component view.

> ***Warning***
>
> *Your application will be offline from now until you complete the process in step 9*

6. On the Menu bar, click **Reconfigure**. Note the Template options section with inline help and configuration mode options. A number of additional questions appear if you select Advanced mode.
7. In the **Virtual Server and Pool** section, in the **What FQDNs will clients use to access the servers** question, you must add the host name.
8. No additional changes are necessary, but you may modify any of the other settings as applicable for your implementation. Use the inline help and this deployment guide for information on specific settings.
9. Click **Finished**. The upgrade is now complete and all applicable objects appear in the Component view.

Troubleshooting

Symptom – Microsoft Windows Server 2008 or 2012 server receives a TCP SYN request from the client, but when using nPath, the server is not sending a SYN ACK response.

Resolution – If you are having issues using direct server return (aka nPath) with a Microsoft Windows Server 2008 or 2012, make sure you alias the nPath virtual server using a loopback adapter and properly alias the servers local adaptors.

Use the following guidance to properly alias server local adaptors on Microsoft servers:

1. Install the loopback adapter as it is not installed by default.
2. From the command line, issue the command **ipconfig /all** to determine the servers local interface names. Locate the loopback adapter information in the results. For example:

   ```
   Ethernet adapter Local Area Connection 2:
   Description ...Microsoft Loopback Adapter
   ```

3. Alias the virtual server IP to the loopback using the following netsh command

   ```
   netsh interface ipv4 add address "Local Area Connection 2" vs_ip_address subnet_mask
   ```

4. Enable forwarding for all of the interfaces on the machine. For example:

   ```
   netsh interface ipv4 set interface "Local Area Connection" forwarding=enabled
   netsh interface ipv4 set interface "Local Area Connection 2" forwarding=enabled
   ```

   You should receive an "OK" message after each command has been executed.

5. The TCP/IP stack in Windows Server 2008/2012 supports strong host sends and receives for both IPv4 and IPv6 by default. Weak send and receive should be enabled explicitly using following commands:

   ```
   netsh interface ipv4 (or ipv6) set interface "loopback interface name" weakhostreceive=enabled
   netsh interface ipv4 (or ipv6) set interface "loopback interface name" weakhostsend=enabled
   netsh interface ipv4 set interface "local interface name" weakhostreceive=enabled
   ```

Appendix: Manual configuration table

We strongly recommend using the iApp template to configure the BIG-IP system for nPath. This table contains a list of BIG-IP configuration objects along with any non-default settings you should configure as a part of this deployment. Unless otherwise specified, settings not mentioned in the table can be configured as applicable for your configuration. For specific instructions on configuring individual objects, see the online help or product manuals.

BIG-IP LTM Object	Non-default settings/Notes		
Health Monitor (*Main tab-->Local Traffic -->Monitors*)	***Name***	Type a unique name	
	Type	**ICMP, TCP** or **UDP** depending on your configuration	
	Interval	**30** (recommended)	
	Timeout	**91** (recommended)	
Pool (*Main tab-->Local Traffic -->Pools*)	***Name***	Type a unique name	
	Health Monitor	Select the monitor you created above	
	Slow Ramp Time[1]	**300**	
	Load Balancing Method	Choose a load balancing method. We use the default **Least Connections (member)**	
	Address	Type the IP Address of the servers	
	Service Port	Type the appropriate port. Click **Add** to repeat Address and Service Port for all servers.	
Profiles (*Main tab-->Local Traffic -->Profiles*)	***Persistence*** *(optional)* (*Profiles-->Persistence)*	Name	Type a unique name
		Persistence Type	if your deployment requires persistence, select the appropriate method.
	Fast L4 (*Profiles-->Protocol*)	Name	Type a unique name
		Parent Profile	**Fast L4**
Virtual Server (*Main tab-->Local Traffic -->Virtual Servers*)	***Name***	Type a unique name.	
	Type	**Performance (Layer 4)**	
	Address	Type the IP address for the virtual server	
	Service Port	Type the appropriate port.	
	Protocol	Select the appropriate protocol (TCP, UDP, or *All Protocols)	
	Protocol Profile (Client)	Select the Fast L4 profile you created above	
	Address Translation	Clear the box to disable Address translation. nPath does not function properly if this is enabled.	
	Port Translation	Clear the box to disable Port translation. nPath does not function properly if this is enabled.	
	iRule	If applicable, enable any iRules you created	
	Default Pool	Select the pool you created above	
	Persistence Profile	If applicable, select the Persistence profile you created	

[1] You must select Advanced from the Configuration list for these options to appear

Glossary

application service
iApps application services use an *iApp Template* to guide users through configuring new BIG-IP® system configurations. An application service lets an authorized user easily and consistently deploy complex BIG-IP® system configurations just by completing the information required by the associated template. Every application service is attached to a specific configuration and cannot be copied the way that iApps templates can.

iApp Template
iApps templates create configuration-specific forms used by application services to guide authorized users through complex system configurations. The templates provide programmatic, visual layout and help information. Each new application service uses one of the templates to create a screen with fields and help that guide the user through the configuration process and creates the configuration when finished.

iApps templates allow users to customize by either modifying an existing template or creating one from scratch. Users can create scratch-built templates using either the iApps Templates screen or any text-editing software.

configuration utility
The Configuration utility is the browser-based application that you use to configure the BIG-IP system.

custom profile
A custom *profile* is a profile that you create. A custom profile can inherit its default settings from a parent profile that you specify. See also parent profile.

health monitor
A health monitor checks a node to see if it is up and functioning for a given service. If the node fails the check, it is marked down. Different monitors exist for checking different services.

iRule
An iRule is a user-written script that controls the behavior of a connection passing through the BIG-IP system. iRules™ are an F5 Networks feature and are frequently used to direct certain connections to a non-default load balancing pool. However, iRules can perform other tasks, such as implementing secure network address translation and enabling session persistence. You can attach iRules you created to your nPath routing application service in the advanced configuration mode.

load balancing method
A load balancing method or algorithm is a particular method of determining how to distribute connections across a *load balancing pool*. There are several different load balancing methods on the BIG-IP system. If you are working with servers that differ significantly in processing speed and memory, you might want to use a method such as Ratio or Weighted Least Connections.

Load balancing calculations can be localized to each pool (member-based calculation) or they may apply to all pools of which a server is a member (node-based calculation). For detailed information, see the product documentation.

See the table on the following page for a description of most load balancing methods.

Method	Description	When to use
Round Robin	Round Robin mode passes each new connection request to the next server in line, eventually distributing connections evenly across the array of machines being load balanced.	Round Robin mode works well in most configurations, especially if the equipment that you are load balancing is roughly equal in processing speed and memory.
Ratio (member) **Ratio (node)**	The LTM distributes connections among pool members in a static rotation according to ratio weights you define. The number of connections each system receives over time is proportionate to the ratio weight you defined for each pool member. You set a ratio weight when you add each pool member in the iApp.	These are static load balancing methods, basing distribution on user-specified ratio weights that are proportional to the capacity of the servers.

Method	Description	When to use
Dynamic Ratio (member) **Dynamic Ratio (node)**	The Dynamic Ratio methods select a server based on various aspects of real-time server performance analysis. These methods are similar to the Ratio methods, except the ratio weights are system-generated, and the values of the ratio weights are not static. These methods are based on continuous monitoring of the servers, and the ratio weights are therefore continually changing.	The Dynamic Ratio methods are used specifically for load balancing traffic to RealNetworks® RealSystem® Server platforms, Windows® platforms equipped with Windows Management Instrumentation (WMI), or any server equipped with an SNMP agent such as the UC Davis SNMP agent or Windows 2000 Server SNMP agent. Note: To implement Dynamic Ratio load balancing, you must first install and configure the necessary server software for these systems, and then install the appropriate performance monitor.
Fastest (node) **Fastest (application)**	The Fastest load balancing mode load balances based upon the number of outstanding Layer 7 requests to a pool member and the number of open L4 connections.	The Fastest methods are useful in environments where nodes are distributed across separate logical networks.
Least Connections (member) **Least Connections (node)**	The Least Connections load balancing mode is a dynamic load balancing algorithm that distributes connections to the server that is currently managing the fewest open connections at the time the new connection request is received.	The Least Connections methods function best in environments where the servers have similar capabilities. Otherwise, some amount of latency can occur. If you have servers with varying capacities, consider using the Weighted Least Connections methods instead.
Weighted Least Connections (member) **Weighted Least Connections (node)**	Specifies that the system passes a new connection to the pool member that is handling the lowest percentage of the specified maximum number of concurrent connections allowed. This mode requires that you specify a value for the connection-limit option for all members of the pool.	This mode works best in environments where the servers or other equipment you are load balancing have different but quantified capability limits.
Observed (member) **Observed (node)**	With the Observed methods, nodes are ranked based on the number of connections. The Observed methods track the number of Layer 4 connections to each node over time and create a ratio for load balancing	The need for the Observed methods is rare, and they are not recommended for large pools.
Predictive (member) **Predictive (node)**	The Predictive methods use the ranking methods used by the Observed methods. However, with the Predictive methods, LTM analyzes the trend of the ranking over time, determining whether a nodes performance is currently improving or declining. The servers with performance rankings that are currently improving receive a higher proportion of the connections.	The need for the Predictive methods is rare, and they are not recommended for large pools.
Least Sessions	The Least Sessions method selects the server that currently has the least number of entries in the persistence table. Use of this load balancing method requires that the virtual server reference a type of profile that tracks persistence connections, such as the Source Address Affinity or Universal profile type. Note: The Least Sessions methods are incompatible with cookie persistence.	The Least Sessions method works best in environments where the servers or other equipment that you are load balancing have similar capabilities.

load balancing pool

A load balancing pool is a logical set of devices, such as Web servers, that you group together to receive and process traffic. Instead of sending client traffic to the destination IP address specified in the client request, Local Traffic Manager sends the request to any of the servers that are members of that pool. This helps to efficiently distribute the load on your server resources.

profile

Profiles are a configuration tool that you can use to affect the behavior of certain types of network traffic. More specifically, a profile is an object that contains settings with values, for controlling the behavior of a particular type of network traffic. Profiles also provide a way for you to enable connection and session persistence, and to manage client application authentication.

self IP address

Self IP addresses are the IP addresses owned by the BIG-IP system that you use to access the internal and external VLANs.

SNAT
A SNAT (Secure Network Address Translation) is a feature that defines a routable alias IP address that one or more nodes can use as a source IP address when making connections to hosts on the external network.

SNAT pool
A SNAT pool is a pool of translation addresses that you can map to one or more original IP addresses. Translation addresses in a SNAT pool are not self IP addresses.

virtual server
A virtual server is a traffic-management object on the BIG-IP system that is represented by an IP address and a service port. This is the address clients use to connect to the servers (or a FQDN resolves to this address). The BIG-IP intercepts the client request, and then directs the traffic according to your configuration instructions.

VLAN
A VLAN is a logical grouping of interfaces connected to network devices. You can use a VLAN to logically group devices that are on different network segments. Devices within a VLAN use Layer 2 networking to communicate and define a broadcast domain.

Document Revision History

Version	Description	Date
1.0	New Deployment Guide for BIG-IP v11.4	06-11-2013
1.1	- Added support for BIG-IP v11.4.1 and 11.5. - Add the section *Troubleshooting on page 13* with an entry about Windows Server 2008 or 2012. - The updated iApp template in 11.5 no longer enables Source and Port translation on the virtual server. Made a note in the manual configuration table not to enable Source or Port translation for nPath deployments.	01-31-2014
1.2	Added support for BIG-IP v11.5.1 and 11.6.	08-25-2014

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com