Deploying the BIG-IP System with HTTP Applications

Welcome to the F5® deployment guide for HTTP applications. This document contains guidance on configuring the BIG-IP® system version 11.4 and later for most web server implementations, resulting in a secure, fast, and available deployment. This guide shows how to quickly and easily configure the BIG-IP system using the HTTP iApp Application template. There is also an appendix with manual configuration tables for users who prefer to create each individual object.

Why F5?

The BIG-IP system provides a number of ways to accelerate, optimize, and scale HTTP deployments. When the BIG-IP system relieves web servers from tasks such as compression, caching, and SSL processing, each server is able to devote more resources to running applications and can service more user requests.

Products and applicable versions

Product	Versions
BIG-IP LTM, AAM, AFM, ASM	11.4 - 13.0 (ASM available in 12.0 and later only)
HTTP applications	Not applicable
HTTP iApp template	System iApp that ships with v11.4 and later
Deployment Guide version	2.8 (see *Document Revision History on page 38*)
Last updated	05-12-2017

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

If you are looking for older versions of this or other deployment guides, check the Deployment Guide Archive tab at: https://f5.com/solutions/deployment-guides/archive-608

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

Contents

What is F5 iApp?

Introduced in version 11 of the BIG-IP system, 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 HTTP applications acts as the single-point interface for building, managing, and monitoring these servers. For more information on iApp, see the White Paper *F5 iApp: Moving Application Delivery Beyond the Network at 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 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 version, and have an existing Application Service that used the f5.http iApp template, see *Upgrading an Application Service from previous version of the iApp template on page 26*.

➤ This document provides guidance for using the iApp for HTTP applications 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, because the configuration can be complex, we recommend using the iApp template.

➤ If you are using the BIG-IP system to offload SSL or for SSL Bridging, we assume you have already obtained the appropriate SSL certificate and key, and it is installed on the BIG-IP LTM system.

➤ ***NEW:*** The *downloadable version of the HTTP iApp is now fully supported by F5 Networks*. For information on how to deploy the downloadable version, see *http://f5.com/pdf/deployment-guides/esd-iapp-http-dg.pdf.*

  • If you want to use BIG-IP Access Policy Manager for your HTTP application, you must use the downloadable iApp.

➤ If you are deploying a specific application, we recommend you first check the deployment guide index at *https://f5.com/solutions/deployment-guides* to see if there is a deployment guide for your application.

➤ If you plan to use the ASM module (v12.0+ only), you must have created a LTM Policy object with ASM enabled.

➤ If you are using the BIG-IP Application Acceleration Manager (AAM) for Symmetric optimization between two BIG-IP systems (optional), you must have pre-configured the BIG-IP AAM for Symmetric Optimization using the Quick Start wizard or manually configured the necessary objects. See the AAM documentation (*http://support.f5.com/kb/en-us/products/big-ip-aam.html*) for specific instructions on configuring BIG-IP AAM for Symmetric Optimization.

Skip ahead   `Advanced`

If you are familiar with the HTTP iApp or the BIG-IP system, you can skip the Configuration Scenario and Preparation sections. See:

  • *Configuring the BIG-IP iApp for HTTP applications on page 8* if using the iApp template, or
  • *Appendix: Manual configuration table on page 28* if configuring the BIG-IP system manually.

Optional modules

This HTTP iApp allows you to use the following optional modules on the BIG-IP system. To take advantage of these modules, they must be licensed and provisioned before starting the iApp template. For more information on licensing modules, contact your sales representative.

  • **BIG-IP AAM** (formerly BIG-IP WAN Optimization Manager and WebAccelerator)
    BIG-IP AAM provides application, network, and front-end optimizations to ensure consistently fast performance for today's dynamic web applications, mobile devices, and wide area networks. With sophisticated execution of caching, compression, and image optimization, BIG-IP AAM decreases page download times. You also have the option of using BIG-IP AAM for symmetric optimization between two BIG-IP systems. For more information on BIG-IP AAM, see *http://www.f5.com/products/big-ip/big-ip-application-acceleration-manager/overview/*.

  • **Application Visibility and Reporting**
    F5 Analytics (also known as Application Visibility and Reporting or AVR) is a module that lets customers view and analyze metrics gathered about the network and servers as well as the applications themselves. Making this information available from a dashboard-type display, F5 Analytics provides customized diagnostics and reports that can be used to optimize

application performance and to avert potential issues. The tool provides tailored feedback and recommendations for resolving problems. Note that AVR is licensed on all systems, but must be provisioned before beginning the iApp template.

- **BIG-IP AFM**
  BIG-IP Advanced Firewall Manager (AFM) is a high-performance, stateful, full-proxy network firewall designed to guard data centers against incoming threats that enter the network on the most widely deployed protocols—including HTTP/S, SMTP, DNS, and FTP. By aligning firewall policies with the applications they protect, BIG-IP AFM streamlines application deployment, security, and monitoring. For more information on BIG-IP AFM, see *https://f5.com/products/modules/advanced-firewall-manager*.

- **BIG-IP ASM**
  BIG-IP ASM protects the HTTP applications your business relies on with an agile, certified web application firewall and comprehensive, policy-based web application security. Offering threat assessment and mitigation, visibility, and almost limitless flexibility, BIG-IP ASM helps you secure your HTTP applications. For more information on BIG-IP Application Security Manager, see *http://www.f5.com/products/big-ip/big-ip-application-security-manager/overview/*.

Configuration scenarios

With the iApp template for HTTP servers, you can configure the BIG-IP system to optimize and direct traffic to HTTP servers with ease. You can also configure the BIG-IP system for different system scenarios using the options found in the iApp, as described in this section.

Configuring the BIG-IP system as reverse (or inbound) proxy

In its traditional role, the BIG-IP system is a reverse proxy. The system is placed in the network between the clients and the servers. Incoming requests are handled by the BIG-IP, which interacts on behalf of the client with the desired server or service on the server. This allows the BIG-IP system to provide scalability, availability, server offload, and much more, all completely transparent to the client.


Figure 1: *Using the BIG-IP system as a reverse proxy*

To configure this scenario

There are no questions in the iApp template that you must answer in a specific way for the BIG-IP system to act as a reverse proxy, the BIG-IP system acts as a reverse proxy by default.

Accelerating application traffic over the WAN

The iApp enables you to use the BIG-IP system's Application Acceleration Manager module to optimize and secure your web traffic over the WAN (wide area network). The iApp uses the default *iSession profile* to create a secure tunnel between BIG-IP systems to accelerate and optimize the traffic.

In this scenario, you must have a symmetric BIG-IP deployment (as shown in Figure 2), with a BIG-IP system between your clients and the WAN, and another between the WAN and your web servers. You run the iApp template on each of the BIG-IP systems, using the settings found in the following table.


Figure 2: *Using an iSession tunnel to secure and optimize traffic between two BIG-IP systems*

To configure this scenario

If you select this option, you must have already configured the BIG-IP AAM for Symmetric Optimization as mentioned in the prerequisites. See the BIG-IP AAM documentation available on AskF5™ (*http://support.f5.com/kb/en-us/products/big-ip-aam.html*) for specific instructions on configuring BIG-IP AAM for Symmetric Optimization.

To configure the system for this scenario, at a minimum you must answer the following questions with the appropriate answers in the iApp template as shown in the following table.

The table assumes you are configuring the BIG-IP system on the client side of the WAN.

iApp template question	Your answer
On the BIG-IP system between <u>clients</u> and the WAN	
What type of network connects clients to the BIG-IP system? (on page 9)	**LAN** or **WAN** as appropriate
What type of network connects servers to the BIG-IP system? (on page 10)	**WAN through another BIG-IP system**
Do you want to create a new pool or use an existing one?	Typically you would leave this at the default for this scenario (Do not use a pool), however you could create a pool of local servers to use as a fallback in case the WAN becomes unavailable.
On the BIG-IP system between <u>servers</u> and the WAN	
What type of network connects clients to the BIG-IP system? (on page 9)	**WAN through another BIG-IP system**
What type of network connects servers to the BIG-IP system? (on page 10)	**LAN** or **WAN** as appropriate (Typically LAN)

Using the BIG-IP system with SSL traffic

The HTTP iApp template provides the following options for dealing with encrypted traffic:

- **SSL Offload**
  When performing SSL offload, the BIG-IP system accepts incoming encrypted traffic, decrypts (or terminates) it, and then sends the traffic to the servers unencrypted. By saving the servers from having to perform the decryption duties, F5 improves server efficiency and frees server resources for other tasks. SSL certificates and keys are stored on the BIG-IP system.

- **SSL Bridging**
  With SSL Bridging, also known as SSL re-encryption, the BIG-IP system accepts incoming encrypted traffic, decrypts it for processing, and then re-encrypts the traffic before sending it back to the servers. This is useful for organizations that have requirements for the entire transaction to be SSL encrypted. In this case, SSL certificates and keys must be are stored and maintained on the BIG-IP system and the web servers.

- **SSL pass-through**
  With SSL pass-through, the BIG-IP system does not process the encrypted traffic at all, just sends it on to the servers.

- **No SSL (plaintext)**
  In this scenario, the BIG-IP system does not perform any SSL processing, as all traffic is only plaintext.

- **Server-side encryption**
  In this scenario, the BIG-IP system accepts unencrypted traffic and then encrypts is before sending it to the servers. While more uncommon than offload or bridging, this can be useful for organizations that require all traffic behind the system to be encrypted.
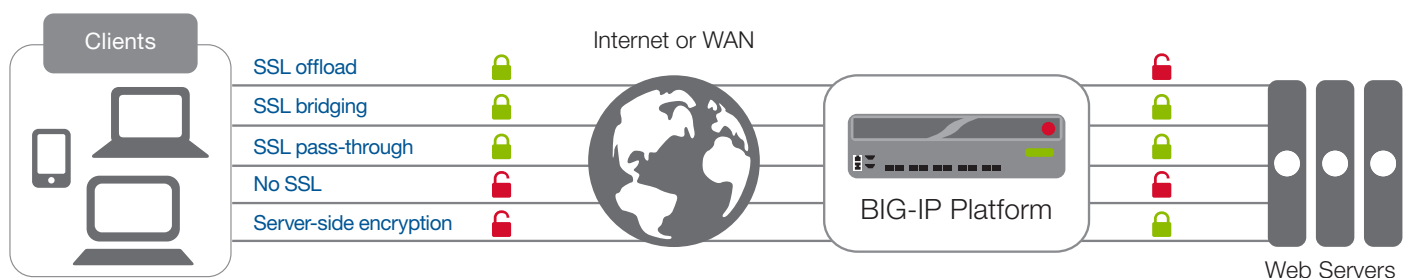
Figure 3: *SSL options*

To configure these scenarios

For SSL offload or SSL bridging, you must have imported a valid SSL certificate and key onto the BIG-IP system. Importing certificates and keys is not a part of the template, see **System > File Management > SSL Certificate List**, and then click **Import**.

iApp template question	Your answer
How should the BIG-IP system handle SSL traffic (on page 11)	Select the appropriate option for your configuration: SSL Offload: **Encrypt to clients, plaintext to servers** SSL Bridging: **Terminate SSL from clients, re-encrypt to servers** SSL Pass-Through: **Encrypted traffic is forwarded without decryption** No SSL: **Plaintext to clients and servers** Server-side encryption: **Plaintext to clients, encrypt to servers**

Using this guide

This deployment guide is intended to help users deploy web-based applications using the BIG-IP system. This document contains guidance configuring the BIG-IP system using the iApp template, as well as manually configuring the BIG-IP system.

Using this guide to configure the iApp template

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

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 UI for the iApp template itself are all displayed 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 the drop-down list)
   - **Choice #2** (in the drop-down list)
     a. *Second level question dependent on selecting choice #2*
        - **Sub choice #1**
        - **Sub choice #2**
          a. *Third level question dependent on sub choice #2*
             - **Sub-sub choice**
             - **Sub-sub #2**
               a. *Fourth level question (and so on)*

Advanced options and questions in the template are marked with the Advanced icon: **Advanced** . In the iApp template, these questions only appear if you select the Advanced configuration mode.

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 HTTP 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 28.*

Preparing to use the iApp

In order to use the iApp for HTTP applications, 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.

BIG-IP System Preparation Table		
Basic/Advanced mode	In the iApp, you can configure the system for your HTTP application with F5 recommended settings (Basic mode) which are a result of extensive testing and tuning with a wide variety of HTTP applications. Advanced mode allows configuring the BIG-IP system on a much more granular level, configuring specific options, or using your own pre-built profiles or iRules. Basic/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 8)	
Network	**Type of network between <u>clients</u> and the BIG-IP system**	**Type of network between <u>servers</u> and the BIG-IP system**
	LAN | WAN | WAN through another BIG-IP system	LAN | WAN | WAN through another BIG-IP system
	If WAN through another BIG-IP system, you must have BIG-IP AAM pre-configured for Symmetric Optimization.	
	Where are BIG-IP virtual servers in relation to the servers	**Expected number of concurrent connections per server**
	Same subnet | Different subnet	More than 64k concurrent | Fewer than 64k concurrent
	If they are on different subnets, you need to know if the web servers have a route through the BIG-IP system. If there is not a route, you need to know the number of concurrent connections.	If more than 64k per server, you need an available IP address for each 64k connections you expect for the SNAT Pool
SSL Encryption	**SSL offload or SSL bridging**	**Re-encryption (Bridging and server-side encryption)**
	If configuring the system for SSL Offload or SSL Bridging, you must have imported a valid SSL certificate and key onto the BIG-IP system. You have the option of also using an Intermediate (chain) certificate as well if required in your implementation. *Certificate*: *Key*: *Intermediate Certificate (optional)*:	When the BIG-IP system encrypts traffic to the servers, it is acting as an SSL client and by default we assume the servers do not expect the system to present its client certificate on behalf of clients traversing the virtual server. If your servers expect the BIG-IP system to present a client certificate, you must create a custom Server SSL profile outside of the template with the appropriate certificate and key.
Virtual Server and Pools	**Virtual server**	**HTTP server pool**
	The virtual server is the address clients use to access the servers.	The *load balancing pool* is the LTM object that contains the servers.
	IP address for the virtual server: *Associated service port*: *FQDN clients will use to access the HTTP servers*:	*IP addresses of the servers:* 1: 2: 3: 4: 5: 6: 7: 8: 9:
Profiles	The iApp template can create *profiles* using the F5 recommended settings, or you can choose **Do not use** many of these profiles). F5 recommends using the profiles created by the iApp; however you also 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 the following profiles (some only in Advanced mode). Any profiles must be present on the system before you can select them in the iApp.	
	HTTP | Persistence | HTTP Compression | TCP LAN | TCP WAN | OneConnect | Web Acceleration | NTLM | iSession	
Health monitor	**HTTP request**	**User account**
	In Advanced mode, you have the option of selecting the type of HTTP request the health monitor uses: GET or POST. You can also specify Send and Receive strings to more accurately determine server health. *Send string* (the URI sent to the servers): *Receive string (*what the system expects in return): *POST Body* (only if using POST):	Also in advanced mode, the monitor can attempt to authenticate to the web servers as a part of the health check. If you want the monitor to require credentials, create a user account specifically for this monitor that has no additional permissions and is set to never expire. Account maintenance becomes a part of the health monitor, as if the account is deleted or otherwise changed, the monitor will fail and the servers will be marked down.
Application Acceleration Manager	You can optionally use the BIG-IP Application Acceleration Manager (AAM) module to help accelerate your HTTP traffic. To use BIG-IP AAM, it must be fully licensed and provisioned on your BIG-IP system. Consult your F5 sales representative for details. If you are using BIG-IP AAM, and want to use a custom Web Acceleration policy, it must have an Acceleration policy attached.	
iRules	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.	

Configuring the BIG-IP iApp for HTTP applications

Use the following guidance to configure the BIG-IP system for HTTP applications using the HTTP iApp template.

Getting Started with the iApp for HTTP applications

To begin the HTTP iApp Template, use the following procedure. If you want to use the supported downloadable HTTP iApp, see *http://f5.com/pdf/deployment-guides/esd-iapp-http-dg.pdf*. This version has features not available in the base f5.http iApp.

1. Log on to the BIG-IP system.

2. On the Main tab, expand **iApps**, 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 **HTTP-iapp_.**

5. From the **Template** list, select **f5.http**. The HTTP 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, introduced in 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.

     As mentioned, 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 the Advanced icon.

Network

This section contains questions about your networking configuration.

1. *What type of network connects clients to the BIG-IP system?*
    Choose the type of network that connects your clients to the BIG-IP system. If you choose WAN or LAN, the BIG-IP system uses this information to determine the appropriate TCP optimizations. If you choose WAN through another BIG-IP system, the system uses a secure an optimized tunnel (called an iSession tunnel) for traffic between BIG-IP systems on opposite sides of the WAN. Only choose this option if you have another BIG-IP system across the WAN that will be a part of this implementation.

    *   **Local area network (LAN)**
        Select this option if most clients are connecting to the BIG-IP system on a LAN. This field is used to determine the appropriate TCP profile which is optimized for LAN clients. In this case, the iApp creates a TCP profile using the *tcp-lan-optimized* parent with no additional modifications.

    *   **Wide area network**
        Select this option if most clients are connecting to the BIG-IP system over a WAN. This field is used to determine the appropriate TCP profile which is optimized for WAN clients. In this case, the iApp creates a TCP profile using the *tcp-wan-optimized* parent with no additional modifications.

    *   **WAN through another BIG-IP system**
        Select this option if client traffic is coming to this BIG-IP system from a remote BIG-IP system across a WAN. As mentioned in the introduction to this question, the iApp creates an iSession tunnel between this BIG-IP system and the BIG-IP system you will configure (or already have configured) on the other side of the WAN.

        If you select this option, you must have already initially configured the BIG-IP AAM for Symmetric Optimization. See the BIG-IP AAM documentation available on Ask F5 for specific instructions on configuring BIG-IP AAM for Symmetric Optimization.

2. *Do you want to restrict client traffic to specific VLANs? (11.5 and later)* `Advanced`
    Which VLANs transport client traffic? (11.4.x) `Advanced`
    The BIG-IP system allows you to restrict client traffic to specific VLANs that are present on the system. This can provide an additional layer of security, as you can allow or deny traffic from the VLANs you choose. By default, all VLANs configured on the system are enabled. If you select to enable or disable traffic on specific VLANs, you must specify the VLANs in the next question. The VLAN objects must already be configured on this BIG-IP system before you can select them.

    In version 11.4.x, you can only allow traffic from specific VLANs using the iApp template; v11.5 and later enables you to allow or deny client traffic from specific VLANs. If using v11.4.x, all allowed VLANs appear in the Selected list. Use the Move buttons (<<) and (>>) to adjust list membership. Only VLANs in the Selected list are allowed. With 11.4.x, you do NOT see the following options.

    *   **Enable traffic on all VLANs and Tunnels**
        Choose this option to allow traffic from all VLANs and Tunnels. If you select this option, the question asking about VLANs disappears. Continue with #3.

    *   **Yes, enable traffic only on the VLANs I specify**
        Choose this option to restrict client traffic to specific VLANs that you choose in the following question. The system will accept client traffic from these VLANs, and deny traffic from all other VLANs on the system.

        a. *On which VLANs should traffic be enabled or disabled?*
            Use this section to specify the VLANs that will accept client traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so click the VLANs and then use the Move buttons to adjust list membership.

            > ➡ **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).*

    *   **Yes, disable traffic only on the VLANs I specify**
        Choose this option to deny client traffic from the specific VLANs that you choose in the following question. The system will refuse client traffic from these VLANs, and accept traffic from all other VLANs on the system.

        a. *On which VLANs should traffic be enabled or disabled?*
            Use this section to specify the VLANs that should not accept client traffic. By default, all VLANs on the BIG-IP system appear in the Selected box, so it is critical in this case that you click the VLANs and then use the Move button (>>) to adjust list membership.

> ⚠️ **Warning** *If you choose to disable certain VLANs, you must move at least one VLAN to the Options list. Otherwise, the system will deny traffic from all VLANs on the box, and the configuration, although valid, will not pass any traffic.*

3. *What type of network connects servers to the BIG-IP system?*
   Choose the type of network that connects your servers to the BIG-IP system. Similar to the question about clients connecting to the BIG-IP system, if you choose WAN or LAN, the system uses this information to determine the appropriate TCP optimizations. If you choose WAN through another BIG-IP system, the system uses a secure an optimized tunnel (called an iSession tunnel) for traffic between BIG-IP systems on opposite sides of the WAN. Only choose this option if you have another BIG-IP system across the WAN that will be a part of this HTTP implementation.

   • **Local area network (LAN)**
     Select this option if the servers connect to the BIG-IP system on a LAN. This field is used to determine the appropriate TCP profile. In this case, the iApp creates a TCP profile using the *tcp-lan-optimized* parent with no additional modifications.

   • **Wide area network**
     Select this option if the servers connect to the BIG-IP system over a WAN. This field is used to determine the appropriate TCP profile. In this case, the iApp creates a TCP profile using the *tcp-wan-optimized* parent with no additional modifications.

   • **WAN through another BIG-IP system**
     Select this option if servers are across a WAN behind another BIG-IP system. As mentioned in the introduction to this question, the iApp creates an iSession tunnel between this BIG-IP system and the BIG-IP system you will configure (or already have configured) on the other side of the WAN.

     If you select this option, you must have already initially configured the BIG-IP AAM for Symmetric Optimization. See the BIG-IP AAM documentation available on Ask F5 for specific instructions on configuring BIG-IP AAM for Symmetric Optimization.

4. *Where will your BIG-IP virtual servers be in relation to the web servers?*
   Select whether your BIG-IP virtual servers are on the same subnet as your web servers, or on different subnets. This setting is used to determine the secure NAT (*SNAT*) and routing configuration.

   • **BIG-IP virtual server IP and web servers are on the same subnet**
     If the BIG-IP virtual servers and HTTP servers are on the same subnet, SNAT is configured on the BIG-IP virtual server and you must specify the number of concurrent connections.

     a. *How many connections to you expect to each web server?*
        Select whether you expect more or fewer than 64,000 concurrent connections to each server. This answer is used to determine what type of SNAT that system uses. A SNAT is an object that maps the source client IP address in a request to a translation address defined on the BIG-IP device. The system configures SNAT Auto Map (fewer than 64,000) or a SNAT pool (more than 64,000).

        • **Fewer than 64,000 concurrent connections**
          Select this option if you expect fewer than 64,000 concurrent connections per web server. With this option, the system applies SNAT Auto Map, doesn't require any additional IP addresses, as an existing self IP address is used for translation. Continue with *Virtual Server and Pools on page 15.*

        • **More than 64,000 concurrent connections**
          Select this option if you have a very large deployment and expect more than 64,000 connections at one time. The iApp creates a SNAT Pool, for which you need one IP address for each 64,000 connections you expect.

          a. *Create a new SNAT pool or use an existing one?*
             If you have already created a SNAT pool on the BIG-IP system, you can select it from the list. Otherwise, the system creates a new SNAT pool with the addresses you specify.

             • **Create a new SNAT pool**
               Select this option (the default) to enable the system to create a new SNAT pool. You must specify the appropriate number of IP addresses in the next question.

               a. *Which IP addresses do you want to use for the SNAT pool?*
                  Specify one otherwise unused IP address for every 64,000 concurrent connections you expect, or fraction thereof. Click **Add** for additional rows. Do not use any self IP addresses on the BIG-IP system.

             • **Select a SNAT pool**
               Select the SNAT pool you created for this deployment from the list.

> ⓘ **Important** *If you choose more than 64,000 connections, but do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per server is reached, new requests fail.*

- **BIG-IP virtual servers and web servers are on different subnets**
  If the BIG-IP virtual servers and servers are on different subnets, the following question appears.

  a. *How have you configured routing on your web servers?*
     If you chose different subnets, this question appears asking whether the web servers use this BIG-IP system's self IP address as their default gateway. Select the appropriate answer.

     - **Servers have a route to clients through the BIG-IP system**
       Choose this option if the servers use the BIG-IP system as their default gateway. In this case, no configuration is needed to support your environment to ensure correct server response handling. Continue with the next section.

     - **Servers do not have a route to clients through the BIG-IP system**
       If the web servers do not use the BIG-IP system as their default gateway, *SNAT* is configured on the BIG-IP virtual server and you must select the expected number of concurrent connections in the next question.

       a. *How many connections to you expect to each web server?*
          Select whether you expect more or fewer than 64,000 concurrent connections to each server. This answer is used to determine what type of SNAT that system uses. A SNAT is an object that maps the source client IP address in a request to a translation address defined on the BIG-IP device. The system configures SNAT Auto Map (fewer than 64,000) or a SNAT pool (more than 64,000).

          - **Fewer than 64,000 concurrent connections**
            Select this option if you expect fewer than 64,000 concurrent connections per web server. With this option, the system applies SNAT Auto Map, doesn't require any additional IP addresses, as an existing self IP address is used for translation. Continue with the *SSL Encryption* section.

          - **More than 64,000 concurrent connections**
            Select this option if you have a very large deployment and expect more than 64,000 connections at one time. The iApp creates a SNAT Pool, for which you need one IP address for each 64,000 connections you expect.

            a. *Create a new SNAT pool or use an existing one?*
               If you have already created a SNAT pool on the BIG-IP system, you can select it from the list. Otherwise, the system creates a new SNAT pool with the addresses you specify.

               - **Create a new SNAT pool**
                 Select this option (the default) to enable the system to create a new SNAT pool. You must specify the appropriate number of IP addresses in the next question.

                 a. *Which IP addresses do you want to use for the SNAT pool?*
                    Specify one otherwise unused IP address for every 64,000 concurrent connections you expect, or fraction thereof. Click **Add** for additional rows. Do not use any self IP addresses on the BIG-IP system.

               - **Select a SNAT pool**
                 Select the SNAT pool you created for this deployment from the list.

                 > ⓘ **Important** *If you choose more than 64,000 connections, but do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per server is reached, new requests fail.*

SSL Encryption

Before running the template you should have already imported a certificate and key onto the BIG-IP system. While the BIG-IP system does include a self-signed SSL certificate that can be used internally or for testing, we strongly recommend importing a certificate and key issued from a trusted Certificate Authority for processing client-side SSL.

For information on SSL certificates on the BIG-IP system, see the online help or the *Managing SSL Certificates for Local Traffic* chapter in the **Configuration Guide for BIG-IP Local Traffic Manager** *available at* http://support.f5.com/kb/en-us.html.

1. *How should the BIG-IP system handle SSL traffic?*
   There are four options for configuring the BIG-IP system for SSL traffic. Select the appropriate mode for your configuration.

- **Encrypt to clients, plain text to servers (SSL Offload)**
  Choose this method if you want the BIG-IP system to offload SSL processing from the servers. You need a valid SSL certificate and key for this method.

  a. _Which Client SSL profile do you want to use?_ `Advanced`
     Select whether you want the iApp to create a new Client SSL _profile_, or if you have already created a Client SSL profile which contains the appropriate SSL certificate and key.

     Unless you have requirements for configuring specific Client SSL settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic** > **Profiles** > **SSL** > **Client** to create a Client SSL profile. To select any new profiles you create, you need to restart or reconfigure this template.

     - _Select an existing Client SSL profile_
       If you created a Client SSL profile for this implementation, select it from the list. If you select an existing Client SSL profile, the rest of the questions in this section disappear. Continue with _Virtual Server and Pools on page 15_.

     - **Create a new Client SSL profile**
       Select this option for the iApp to create a new Client SSL profile.

       a. _Which SSL certificate do you want to use?_
          Select the SSL certificate you imported for this implementation.

       b. _Which SSL private key do you want to use?_
          Select the associated SSL private key.

       c. _Which intermediate certificate do you want to use?_ `Advanced`
          If your deployment requires an intermediate or chain certificate, select the appropriate certificate from the list. Immediate certificates are intended to create a chain of trust between the CA that signed the certificate and the CA that is already trusted by the recipient of the certificate. This allows the recipient to verify the validity of the certificates presented, even when the signing CA is unknown.

- **Terminate SSL from clients, re-encrypt to servers (SSL Bridging)**
  Choose this method if you want the BIG-IP system to terminate SSL to process it, and then re-encrypt the traffic to the servers (SSL Bridging). You need a valid SSL certificate and key for the client-side, and optionally for the server-side (see #b).

  a. _Which Client SSL profile do you want to use?_ `Advanced`
     Select whether you want the iApp to create a new Client SSL _profile_, or if you have already created a Client SSL profile which contains the appropriate SSL certificate and key.

     Unless you have requirements for configuring specific Client SSL settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic** > **Profiles** > **SSL** > **Client** to create a Client SSL profile. To select any new profiles you create, you need to restart or reconfigure this template.

     - _Select an existing Client SSL profile_
       If you created a Client SSL profile for this implementation select it from the list. If you select an existing Client SSL profile, the rest of the questions in this section disappear. Continue with _Virtual Server and Pools on page 15_.

     - **Create a new Client SSL profile**
       Select this option for the iApp to create a new Client SSL profile

       a. _Which SSL certificate do you want to use?_
          Select the SSL certificate you imported for this implementation.

       b. _Which SSL private key do you want to use?_
          Select the associated SSL private key.

       c. _Which intermediate certificate do you want to use?_ `Advanced`
          If your implementation requires an intermediate or chain certificate, select the appropriate certificate from the list. Immediate certificates are intended to create a chain of trust between the CA that signed the certificate and the CA that is already trusted by the recipient of the certificate. This allows the recipient to verify the validity of the certificates presented, even when the signing CA is unknown.

  b. _Which Server SSL profile do you want to use?_
     Select whether you want the iApp to create the F5 recommended Server SSL profile, or if you want to choose a Server SSL profile you already created. In this scenario, the BIG-IP system is acting as an SSL client and by default, we assume the

servers do not expect the BIG-IP system to present its client certificate on behalf of clients traversing the virtual server. If your servers expect the BIG-IP system to present a client certificate, you must create a custom Server SSL profile with the appropriate certificate and key.

The default, F5 recommended Server SSL profile uses the serverssl parent profile. For information about the ciphers used in the Server SSL profile, see *http://support.f5.com/kb/en-us/solutions/public/8000/800/sol8802.html*.

- **Encrypted traffic is forwarded without decryption (SSL pass-through)**
  Choose this method if you do not want the BIG-IP system to do anything with encrypted traffic and simply send it to the web servers. This is similar to SSL bridging, although in this case the system does not decrypt then re-encrypt the traffic, it only sends it on to the servers without modification.

  If you select this option, the system changes the default persistence option from Cookie to Source Address Persistence.

- **Plain text to clients, encrypt to servers**
  Choose this method if you want the BIG-IP system to accept plain text from the clients and then encrypt it before sending it to the servers.

  Unless you have requirements for configuring specific Server SSL settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic** > **Profiles** > **SSL** > **Server** to create a Server SSL profile. To select any new profiles you create, you need to restart or reconfigure this template.

  The default, F5 recommended Server SSL profile uses the *serverssl* parent profile. For information about the ciphers used in the Server SSL profile, see *http://support.f5.com/kb/en-us/solutions/public/8000/800/sol8802.html*.

- **Plain text to both clients and servers**
  Choose this method if the BIG-IP system is not sending or receiving any SSL traffic in this implementation.

Application Security Manager (BIG-IP ASM)

The option for deploying BIG-IP ASM only appears in BIG-IP version 12.0 and later.

This section gathers information about BIG-IP Application Security Manager if you want to use it to help protect your deployment. This section only appears if you have fully licensed and provisioned BIG-IP ASM. Contact your F5 sales representative for details.

(i) *Important*   *In order to use BIG-IP ASM, you must have manually created a BIG-IP LTM Policy that includes ASM and applicable Rules. Creating an LTM policy is outside the scope of this guide. If using BIG-IP version 12.1 or 12.1.1, you may need to create the policy from the command line. See the product documentation for specific details. In iApp version 1.2.0rc4 and later, the template can create the ASM policy.*

1. *Do you want to deploy BIG-IP Application Security Manager?*
   Choose whether you want to use BIG-IP ASM, F5's web application firewall to help secure this HTTP deployment. If you choose to use BIG-IP AFM, you must have a BIG-IP LTM policy with ASM enabled.

   - **No, do not use Application Security Manager**
     Select this option if you do not want to enable BIG-IP ASM at this time. You can always re-enter the template at a later date to enable BIG-IP ASM. Continue with the next section.

   - *Select an existing LTM policy from the list*
     If you already created a BIG-IP AFM policy with ASM enabled for this implementation, select it from the list. Only LTM Policy objects with ASM enabled appear in the list.

     a. *Which logging profile would you like to use?*
        The logging profile enables you to log detailed information about BIG-IP ASM events and store those logs on the BIG-IP system or a remote logging server. If you choose to use a logging profile, we strongly recommend creating a custom profile outside the template and then selecting it from the list.

        (i) *Important* *When you store the logs locally, the logging utility may compete for system resources.*

        - **No, do not use a logging profile**
          Select this option if you do not want to use a logging profile at this time. You can always re-enter the template at a later date to enable a logging profile. Continue with the next section.

- *Select an existing logging profile from the list*
  If you already created an ASM logging profile, select it from the list. Only Logging profile objects with ASM enabled appear in the list.

- **Log illegal requests**
  Select this option if you want the system, or remote server, to create a logging profile that only logs illegal requests. By default, the system logs the first 10000 bytes of responses, up to 10 responses per second. You can change the limits by creating a custom profile and using the response logging system variables.

- **Log all requests**
  Select this option if you want the system, or remote server, to log all requests. By default, the system logs the first 10000 bytes of responses, up to 10 responses per second. You can change the limits by creating a custom profile and using the response logging system variables.

Advanced Firewall Manager (BIG-IP AFM)

The option for deploying BIG-IP AFM only appears in BIG-IP version 11.6 and later.

This section gathers information about BIG-IP Advanced Firewall Manager if you want to use it to protect the HTTP deployment. This section only appears if you have fully licensed and provisioned BIG-IP AFM. Contact your F5 sales representative for details.
For information specific on BIG-IP AFM, see *http://support.f5.com/kb/en-us/products/big-ip-afm.html*, and then select your version.

1. *Do you want to use BIG-IP AFM to protect your application?*
   Choose whether you want to use BIG-IP AFM, F5's network firewall, to secure this HTTP deployment. If you choose to use BIG-IP AFM, you can restrict access to the HTTP virtual server to a specific network or IP address. See the BIG-IP AFM documentation for specific details on configuring AFM.

   - **No, do not use Application Firewall Manager**
     Select this option if you do not want to enable BIG-IP AFM at this time. You can always re-enter the template at a later date to enable BIG-IP AFM. Continue with the next section.

   - *Select an existing AFM policy from the list*
     If you already created a BIG-IP AFM policy for this implementation, select it from the list. Continue with **c**.

   - **Yes, use F5's recommended AFM configuration**
     Select this option if you want to enable BIG-IP AFM using F5's recommended configuration.

     a. *Do you want to restrict access to your application by network or IP address?*
        Choose whether you want to restrict access to the HTTP implementation via the BIG-IP virtual server.

        - **No, do not restrict source addresses (allow all sources)**
          By default, the iApp configures the Advanced Firewall module to accept traffic destined for the HTTP virtual server from all sources. If you do not have a need to restrict access to the virtual server, leave this option selected and then continue with **b**.

        - **Restrict source addresses**
          Select this option if you want to restrict access to the HTTP virtual server by IP address or network address.

          a. *What IP or network addresses should be allowed to access your application?*
             Specify the IP address or network access that should be allowed access to the HTTP virtual server. You can specify a single IP address, a list of IP addresses separated by spaces (not commas or other punctuation), a range of IP addresses separated by a dash (for example **192.0.2.10-192.0.2.100**), or a single network address, such as **192.0.2.200/24**.

     b. *How do you want to control access to your application from sources with a low reputation score?*
        The BIG-IP AFM uses an IP intelligence database to categorize IP addresses coming into the system. Choose what you want the system to do for sources that are attempting to access the HTTP virtual server with a low reputation score. For more information, see the BIG-IP AFM documentation.

        (i) *Important   You must have an active IP Intelligence license for this feature to function. See https://f5.com/products/modules/ip-intelligence-services for information.*

        See Troubleshooting on page 27 for a mandatory modification to the configuration if you are using AFM and the IP Intelligence feature to restrict or log traffic with low reputation scores.

- **Allow all sources regardless of reputation**
  Select this option to allow all sources, without taking into consideration the reputation score.

  - **Reject access from sources with a low reputation**
    Select this option to reject access to the HTTP virtual server from any source with a low reputation score.

  - **Allow but log access from sources with a low reputation**
    Select this option to allow access to the HTTP virtual server from sources with a low reputation score, but add an entry for it in the logs.

c. *Would you like to stage a policy for testing purposes?*
   Choose whether you want to stage a firewall policy for testing purposes. A staged policy allows you to evaluate the effect a policy has on traffic by analyzing the system logs, without actually modifying traffic based on the firewall rules. You must already have a policy on the system in order to select it.

   - **Do not apply a staging policy**
     Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.

   - *Select an existing policy from the list*
     If you have already created a firewall policy for this implementation, select it from the list. Only policies that already exist on the system appear in the list. To create a new policy, on the Main tab, click **Security > Network Firewall > Policies**. Specific instructions for creating a firewall policy is outside the scope of this iApp and deployment guide.

d. *Which logging profile would you like to use?*
   Choose whether you or not you want to use a logging profile for this AFM implementation. You can configure the BIG-IP system to log detailed information about BIG-IP system Network Firewall events and store those logs on the BIG-IP system or a remote logging server (supports syslog and Splunk). The list only contains profiles with Network Firewall enabled.
   To create a logging profile, on the Main tab, click **Security > Event Logs > Logging Profiles**. Specific instructions for creating a logging profile is outside the scope of this iApp and deployment guide. See the online help or the *About Local Logging with the Network Firewall* chapter of the **BIG-IP Network Firewall: Policies and Implementations** guide for more information.

   - **Do not apply a logging profile**
     Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.

   - *Select an existing logging profile from the list*
     If you have already created a logging profile for this implementation, select it from the list. You must create a profile before it is available in the list.

Virtual Server and Pools

This section gathers information about your HTTP 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 HTTP deployment via the BIG-IP system.

   If necessary for your configuration, this can be a network address to create a network virtual server (you must specify an IP mask in the following question for a network virtual server). A network virtual server is a virtual server whose IP address has no bits set in the host portion of the IP address (that is, the host portion of its IP address is 0), allowing the BIG-IP system to direct client connections that are destined for an entire range of IP addresses, rather than for a single destination IP address. Thus, when any client connection targets a destination IP address that is in the network specified by the virtual server IP address, the system can direct that connection the pool of HTTP servers.

2. ***If using a network virtual address, what is the IP mask?*** `Advanced`
   If you specified a network address for the virtual server (allowing the virtual server to handle multiple IP addresses), you must enter the full network mask representing the address range. If you specified a single address, you may leave this field blank.

3. ***What port do you want to use for the virtual server?***
   Type the port number you want to use for the BIG-IP virtual server. For HTTP deployments, this is typically 80 (HTTP) or 443 (HTTPS).

4. *Which FQDNs will clients use to access the servers?*
   Type each fully qualified domain name clients will use to access the HTTP deployment. Click the **Add** button to insert additional rows. If you only have one FQDN, do not click Add.

5. *Do you want to redirect inbound HTTP traffic to HTTPS?* `Advanced`
   This question only appears if you selected SSL Offload or SSL Bridging in the SSL question.

   Select whether you want the BIG-IP system to automatically redirect HTTP traffic to the HTTPS virtual server. This is useful when users forget to use HTTPS when attempting to connect to the HTTP deployment.

   • **Redirect HTTP to HTTPS**
     Select this option to redirect HTTP traffic to HTTPS. If you select this option (the default), the BIG-IP system attaches a very small redirect iRule to the virtual server.

     a. *From which port should traffic be redirected?*
        Type the port number for the traffic that you want to redirect to HTTPS. The most common is port 80 (the default).

   • **Do not redirect HTTP to HTTPS**
     Select this option if you do not want to enable the automatic redirect.

6. *Which HTTP profile do you want to use?* `Advanced`
   The HTTP *profile* contains settings for instructing the BIG-IP system how to handle HTTP traffic. Choose whether you want the iApp to create a new HTTP profile or if you have previously created an HTTP profile for this deployment.

   Unless you have requirements for configuring specific HTTP settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic** > **Profiles** > **Services** > **HTTP** to create a HTTP profile. To select any new profiles you create, you need to restart or reconfigure this template.

   • *Select an existing HTTP profile from the list*
     If you already created an HTTP profile for this implementation, select it from the list.

   • **Create a new HTTP profile (recommended)**
     Select this option for the iApp to create a new HTTP profile.

     a. *Should the BIG-IP system insert the X-Forwarded-For header?* `Advanced`
        Select whether you want the BIG-IP system to insert the X-Forwarded-For header in the HTTP header for logging purposes.

        • **Insert the X-Forwarded-For header**
          Select this option if you want the system to include the X-Forwarded-For header. You may have to perform additional configuration on your HTTP servers to log the value of this header. For more information on configuring logging refer to the server documentation.

        • **Do not insert the X-Forwarded-For header**
          Select this option if you do not want the system to include X-Forwarded-For in the HTTP header.

7. *Which persistence profile do you want to use?* `Advanced`
   By using persistence, the system tracks and stores session data, such as the specific pool member that serviced a client request, ensuring client requests are directed to the same pool member throughout the life of a session or during subsequent sessions.

   Unless you have requirements for configuring specific persistence 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** : **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:

   • **Use Cookie Persistence (recommended)** *<if you chose SSL pass-through,"(recommended)" does not appear>*
     Leave this default option to have the BIG-IP system create a new cookie persistence profile (cookie insert mode). With Cookie persistence, the BIG-IP system uses an HTTP cookie stored on the client's computer to allow the client to reconnect to the same server previously visited. We recommend this method for most configurations, except for SSL pass-through.

   • **Source IP Address persistence**
     Select this option if you want to use the Source IP address (also known as simple) persistence. With this mode, the BIG-IP system assigns the built-in Source Address Affinity persistence type, and directs session requests to the same server based only on the source IP address. This is the recommended method if you are using SSL pass-through.

- **Do not use persistence**
  If your implementation does not require persistent connections, select this option.

- *Select an existing persistence profile*
  If you have previously created a persistence profile, you have the option of selecting it instead of allowing the iApp to create a new one. From the list, select an existing persistence profile. We recommend using a persistence profile that uses Cookie persistence, Insert mode.

8. *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 application 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 HTTP 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.

   - **Do not use a pool**
     If you are deploying this iApp in such a way that you do not need a pool of HTTP servers, select this option. If you specified that the servers are connected to the BIG-IP system over the WAN through another BIG-IP system, this is the default option, as the system is sending the traffic across the iSession tunnel to the other BIG-IP system to be distributed to the servers.

   - **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 web server pool. We recommend the default, **Least Connections (member)**.

     b. *Do you want to 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 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 (recommended)**
          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 #c.

          a. *What is the minimum number of active members for each priority 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.

     c. *Which web servers should be included in this pool?*
        Specify the IP address(es) of your web 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.

Delivery Optimization

In this section, you answer questions about how you want the BIG-IP system to optimize the delivery of your HTTP traffic.

1. *Use the BIG-IP Application Acceleration Manager?*
   This question only appears if you have licensed and provisioned the BIG-IP Application Acceleration Manager (AAM).

   Choose whether you want to use the BIG-IP Application Acceleration Manager (formerly known as WebAccelerator). BIG-IP Application Acceleration Manager helps accelerate your HTTP traffic.

- **Yes, use BIG-IP AAM (recommended)**
  Select this option to enable BIG-IP AAM.

- **No, do not use BIG-IP AAM**
  Select this option if you do not want to enable BIG-IP AAM at this time.

2. *Which Web Acceleration profile do you want to use for caching?* `Advanced`
   Select whether you want the system to create a new Web Acceleration profile, or if you have already created a Web Acceleration profile for use in this deployment. The Web Acceleration profile contains the caching settings for this implementation.

   Unless you have requirements for configuring specific acceleration settings (such as specific allowing/denying specific URIs), 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** : **Services : Web Acceleration** to create an acceleration profile. To select any new profiles you create, you need to restart or reconfigure this template.

   Note if using BIG-IP AAM:
   If you are using BIG-IP AAM, and want to select a custom Web Acceleration profile for caching you have already created, it must have an AAM application enabled, otherwise it does not appear in the list of caching profiles. If you want access to all Web Acceleration profiles on the box, then you must choose No to the use BIG-IP AAM question. Use a custom Web Acceleration profile only if you need to define specific URIs that should or should not be cached.

   Note if not using BIG-IP AAM:
   If you are not using BIG-IP AAM, we recommend you only use a custom Web Acceleration profile if you need to define specific URIs which should or should not be cached. You can continue with #6.

   - **Create a profile based on optimized-caching (recommended)**
     Leave this default option to create a new Web Acceleration profile for caching.

   - **Do not use caching**
     This question does not appear if you chose to enable BIG-IP AAM
     Select this option if you do not want to enable caching on the BIG-IP system for this implementation.

   - *Select an existing Web Acceleration profile*
     If you have already created a Web Acceleration profile for your HTTP servers, you can select it from the list.

3. *Do you want to insert the X-WA-Info header?* `Advanced`
   This question only appears if you chose to enable BIG-IP AAM

   The BIG-IP system can optionally insert an X-WA-Info response header that includes specific codes describing the properties and history of the object. The X-WA-Info response header is for informational and debugging purposes only and provides a way for you to assess the effectiveness of your acceleration policy rules.

   By default, the AAM X-WA-info header is not included in the response from the BIG-IP system. If you choose to enable this header, you have two options, Standard and Debug. In Standard mode, the BIG-IP system inserts an HTTP header that includes numeric codes which indicate if and how each object was cached. In Debug mode, the BIG-IP system includes additional information which may help for extended troubleshooting.

   - **Do not insert the header (recommended)**
     Select this option if you do not want to insert the X-WA-Info header. Typically F5 recommends not inserting the header unless instructed to do so by an F5 Technical Support Engineer.

   - **Insert the Standard header**
     Select this option if you want to insert the Standard header. For detailed information on the numeric codes used by the header, see *http://support.f5.com/kb/en-us/solutions/public/13000/700/sol13798.html*

   - **Insert the Debug header**
     Select this option if you want to insert the Debug header for extended troubleshooting.

4. *Do you want to use the legacy AAM performance monitor?* `Advanced`
   This question only appears if you chose to enable BIG-IP AAM

   Enabling the legacy AAM performance monitor can adversely affect system performance. This monitor is primarily used for legacy AAM performance monitoring and debugging purposes, and can adversely affect system performance. The BIG-IP Dashboard provides performance graphs and statistics related to AAM.

- **Do not enable the legacy performance monitor (recommended)**
  Select this option if you do not want to enable the legacy monitor.

- **Enable the legacy performance monitor**
  Select this option if you want to enable the legacy performance monitor. Remember enabling this legacy monitor can impact overall system performance.

  a. _For how many days should the BIG-IP system retain the data?_
     Specify the number of days the BIG-IP system should retain the legacy performance data.

5. _Which acceleration policy do you want to use?_ `Advanced`
   This question only appears if you chose to enable BIG-IP AAM

   Select one of the following predefined acceleration policies from the list.

   - **Generic Policy - Complete**
     This predefined acceleration policy is ideal for Apache HTTP servers, Microsoft Internet Information Services (IIS) web servers, WebLogic application servers, and IBM WebSphere Application Servers. HTML pages are cached and Intelligent Browser Referencing is enabled.

   - **Generic Policy - Enhanced**
     This predefined acceleration policy is ideal for Apache HTTP servers, Internet Information Services (IIS) web servers, WebLogic application servers, and IBM WebSphere Application Servers. HTML pages are cached and Intelligent Browser Referencing is enabled for includes.

   - **Generic Policy - Extension Based.**
     This predefined acceleration policy is ideal for High Performance policy for E-commerce applications that uses File Extensions instead of mime-types. This application policy is ideal if response-based matching is not required.

   - **Generic Policy - Fundamental.**
     This predefined acceleration policy is ideal for Apache HTTP servers, Internet Information Services (IIS) web servers, WebLogic application servers, and IBM WebSphere Application Servers. HTML pages are always proxied and Intelligent Browser Referencing is disabled.

6. _Which compression profile do you want to use?_
   Compression improves performance and end user experience for Web applications that suffer from WAN latency and throughput bottlenecks. Compression reduces the amount of traffic sent to the client to complete a transaction.

   (i) _**Important**_ _If you are using BIG-IP v11.4.x or 11.5.x, see Modifying the configuration produced by the iApp template if using 11.4.x or 11.5.x on page 24 for important information about the F5 recommended HTTP Compression profile._

7. _How do you want to optimize client-side connections?_ `Advanced`
   The client-side TCP profile optimizes the communication between the BIG-IP system and the client by controlling the behavior of the traffic which results in higher transfer rates, improved connection reliability and increased bandwidth efficiency.

   Unless you have requirements for configuring specific TCP optimization 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** > **TCP** to create a TCP profile. To select any new profiles you create, you need to restart or reconfigure this template.

   - **Create the appropriate tcp-optimized profile (recommended)**
     Select this option to have the system create the recommended TCP profile. The parent profile (either WAN or LAN optimized) is determined by your selection to the "What type of network connects clients to the BIG-IP system" question.

   - _Select the TCP profile you created from the list_
     If you created a custom TCP profile for the HTTP servers, select it from the list.

Server offload

In this section, you configure the options for offloading tasks from the servers. This section only appears if you selected Advanced mode.

1. _Which OneConnect profile do you want to use?_ `Advanced`
   OneConnect (connection pooling or multiplexing) improves server scalability by reducing load associated with concurrent

connections and connection rate to HTTP servers. When enabled, the BIG-IP system maintains one connection to each HTTP server which is used to send requests from multiple clients.

Unless you have requirements for configuring specific settings, we recommend allowing the iApp to create a new profile. F5 recommends the default profile which is optimized for most HTTP servers. Creating a custom profile is not a part of this template; see **Local Traffic** >> **Profiles** : **Other : OneConnect** to create a OneConnect profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Create a profile based on the oneconnect parent (recommended)**
  Select this option to have the system create the recommended OneConnect profile. The system uses the oneconnect parent profile with a Source Mask setting of 255.255.255.255.

- **Do not use a OneConnect profile**
  Select this option if you do not require the BIG-IP system to perform connection pooling using a OneConnect profile.

- *Select the OneConnect profile you created from the list*
  If you created a custom OneConnect profile for the HTTP servers, select it from the list.

2. *Which NTLM profile do you want to use?* `Advanced`
   The NTLM profile optimizes network performance when the system is processing NTLM traffic. When both an NTLM profile and a OneConnect profile are enabled, the system can take advantage of server-side connection pooling for NTLM connections.

   If your environment uses NTLM, we recommend allowing the iApp to create a new profile unless you have requirements for configuring specific settings. Creating a custom profile is not a part of this template; see **Local Traffic** >> **Profiles** : **Other : NTLM** to create a NTLM profile. To select any new profiles you create, you need to restart or reconfigure this template.

   - **Use F5's recommended NTLM profile**
     Select this option to have the system create the recommended NTLM profile. The system uses the ntlm parent profile.

   - **Do not use NTLM (recommended)**
     Select this option if you do not use NTLM authentication in your HTTP implementation.

   - *Select the NTLM profile you created from the list*
     If you created a custom NTLM profile for the HTTP servers, select it from the list.

3. *How do you want to optimize server-side connections?* `Advanced`
   The server-side TCP profile optimizes the communication between the BIG-IP system and the server by controlling the behavior of the traffic which results in higher transfer rates, improved connection reliability and increased bandwidth efficiency.

   Unless you have requirements for configuring specific TCP optimization 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 : TCP** to create a TCP profile. To select any new profiles you create, you need to restart or reconfigure this template.

   - **Create the appropriate tcp-optimized profile (recommended)**
     Select this option to have the system create the recommended TCP profile. The parent profile (either WAN or LAN optimized) is determined by your selection to the "What type of network connects servers to the BIG-IP system" question.

   - *Select the TCP profile you created from the list*
     If you created a custom TCP profile for the HTTP servers, select it from the list.

4. *Should the BIG-IP system queue TCP requests?*
   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.

   (i) *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 for the nodes.

- **No, do not enable TCP request queuing (recommended)**
  Select this option if you do not want the BIG-IP system to queue TCP requests.

- **Yes, enable TCP request queuing**
  Select this option if you want to enable TCP request queuing on the BIG-IP system.

  a. _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.

  b. _How many milliseconds should requests remain in the queue?_
     Type a number of milliseconds for the TCP request timeout value.

5. _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 HTTP server over a time period you specify, rather than sending a full proportion of the traffic immediately. Slow Ramp is essential when using the Least Connections load balancing method (our recommended method for web servers), as the BIG-IP 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.

     a. _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 web 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.

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 content that is returned by an HTTP request. The system uses these monitors to ensure traffic is only sent to available HTTP servers.

   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, select it from the list.
     Continue with _iRules on page 23._

   - **Create a new health monitor**
     If you want the iApp to create a new monitor, continue with the following.

     a. _How many seconds should pass between health checks?_ `Advanced`
        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 type of HTTP request should be sent to the servers?_ `Advanced`
        Select whether you want the system to send an HTTP GET or POST request. The GET method requests data from the server, the POST submits data to be processed by the server.

        - **GET**
          Select this option if you want the system to use a GET request. The system uses the URI you specify in the next question to request content from the HTTP server.

- **POST**
  Select this option if you want the system to use a POST request. The system uses the URI you specify in the next question, along with the HTTP POST body you will specify to form the request.

c. *What HTTP URI should be sent to the servers?*
  The HTTP URI is used to specify the resource on the web server for a given request. This parameter can be customized to request a specific part of your application, which can indicate the application-health on a granular level.

d. *What HTTP version do your servers expect clients to use?* `Advanced`
  Choose the HTTP version which you expect most of your clients to be using. This allows the system to detect failures more accurately.

  - **HTTP/1.0**
    Choose this option if you expect your clients to use HTTP/1.0.

  - **HTTP/1.1**
    Choose this option if you expect your clients to use HTTP/1.1.

e. *What HTTP POST body do you want to use for this monitor?* `Advanced`
  This question only appears if you selected a POST request.

  If you selected a POST request, you must specify the message body for the POST.

f. *What is the expected response to the HTTP request?*
  Specify the response you expect returned from the request. The system checks the response from the server against the response you enter here to determine server health.

g. *Should the health monitor require credentials?* `Advanced`
  Choose whether you want the system to attempt to authenticate to the web server deployment as a part of the health check.

  - **No, allow anonymous access**
    Select this option if you do not want the monitor to attempt authentication.

  - **Yes, require credentials for Basic authentication**
    Select this option if you want to attempt Basic authentication as a part of the health monitor. To require credentials, you should have a user account specifically for this health monitor which has no other privileges, and has a password set to never expire.

    a. *What user name should the monitor use?*
      Type the domain and user name for the account you created for the health monitor. You must include the domain in front of the user, such as EXAMPLE\USER.

    b. *What is the associated password?*
      Type the password for the account.

  - **Yes, require credentials for NTLM authentication**
    Select this option if you want to attempt NTLM authentication as a part of the health monitor. To require credentials, you should have a user account specifically for this health monitor which has no other privileges, and has a password set to never expire.

    a. *What user name should the monitor use?*
      Type the user name for the account you created for the health monitor.

      (i) ***Important*** *If you have a complex Active Directory environment where the username is not in the same domain as the Domain Controller being queried, you should prepend the domain name in uppercase letters (DOMAIN\) to the username in this field.*

    b. *What is the associated password?*
      Type the password for the account.

iRules

In this section, you can add custom iRules to the HTTP deployment. This entire section is available only if you selected Advanced mode.

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.

1. ***Do you want to add any custom iRules to the configuration?*** `Advanced`
   Select if have preexisting iRules you want to add to your HTTP implementation.

   ⚠ *Warning   While iRules can provide additional functionality not present in the iApp, improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommended you verify the impact of an iRule prior to deployment in a production environment.*

   If you do not want to add any iRules to the configuration, continue with the following section.

   If you have iRules you want to attach to the virtual server the iApp creates for your web 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.

Statistics and Logging

In this section, you answer questions about optional logging and statistics. This section is available only if you selected Advanced mode.

1. ***Do you want to enable Analytics for application statistics?***
   The Application Visibility Reporting (AVR) module for analytics allows you to view statistics specific to your application implementation. AVR is included and available on all BIG-IP systems v11 and later, however you must have the AVR provisioned for this option to appear. Note that this provisioning requirement is only for AVR, you can view object-level statistics from the BIG-IP system without provisioning AVR.

   ⓘ *Important   Enabling Analytics may adversely affect overall system performance. If you choose to enable Analytics, we recommend gathering statistics for a set time period, such as one week, and then re-entering this template and disabling Analytics while you process the data.*

   If you plan on using AVR for analytics, we recommend creating a custom Analytics profile. To create a new profile, from the Main tab, select **Profiles** and then click **Analytics**. Click **New** and then configure the profile as applicable for your configuration. See the online help or product documentation for specific instructions. To select the new profile, you need to restart or reconfigure the iApp.

   - **Do not enable Application Visibility Reporting**
     If you do not want to enable Analytics, leave this list set to **No**, and continue with the next section.

   - *Select the Analytics profile you created from the list*
     If you choose to enable Analytics, select the Analytics profile you want to use for this implementation from the list.

2. ***Which HTTP request logging profile do you want to use?***
   HTTP request logging enables customizable log messages to be sent to a syslog server for each HTTP request processed by your application. You can choose to enable HTTP request logging by selecting a logging profile you already created from the list. We strongly recommend you thoroughly test the performance impact of using this feature in a staging environment prior to enabling on a production deployment

   Creating a request logging profile is not a part of this template. See Local Traffic>>Profiles: Other: Request Logging. To select any new profiles you create, you need to restart or reconfigure this template.

   - **Do not enable HTTP request logging**
     If you do not want to enable HTTP request logging, leave this list set to **No**, and continue with the next section.

   - *Select the HTTP request logging profile you created from the list*
     If you choose to enable HTTP request logging, select the profile you want to use for this implementation from the list.

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 HTTP application.

Modifying the configuration produced by the iApp template if using 11.4.x or 11.5.x

F5 has discovered that the HTTP compression profile produced by the iApp in versions 11.4 - 11.5 contains an improperly formatted Content Include string, and the omission of this string can lead to poor application performance or unusually high memory consumption on the BIG-IP system.

You must manually create a new HTTP Compression profile and attach it to the virtual server using the iApp, or upgrade to 11.6 or later and update your application service to use the new template.

HTTP Compression Profile (*Local Traffic* > *Profiles* > Services > *HTTP Compression*)	
Name	Type a unique name
Parent Profile	**wan-optimized-compression**
Content List--> Include List *(Copy and paste each entry to the **Content Type** box and click **Include**.)*	text/(css|html|javascript|json|plain|postscript|richtext|rtf|vnd\\.wap\\.wml|vnd\\.wap\\.wmlscript|wap|wml|x-component|x-vcalendar|x-vcard|xml)
	application/(css|css-stylesheet|doc|excel|javascript|json|lotus123|mdb|mpp|ms-excel|ms-powerpoint|ms-word|msaccess|msexcel|mspowerpoint|msproject|msword|photoshop|postscript|powerpoint|ps|psd|quarkexpress|rtf|txt|visio|vnd\\.excel|vnd\\.ms-access|vnd\\.ms-excel|vnd\\.ms-powerpoint|vnd\\.ms-pps|vnd\\.ms-project|vnd\\.ms-word|vnd\\.ms-works|vnd\\.ms-works-db|vnd\\.msaccess|vnd\\.msexcel|vnd\\.mspowerpoint|vnd\\.msword|vnd\\.powerpoint|vnd\\.visio|vnd\\.wap\\.cmlscriptc|vnd\\.wap\\.wmlc|vnd\\.wap\\.xhtml\\+xml|vnd\\.word|vsd|winword|wks|word|x-excel|x-java-jnlp-file|x-javascript|x-json|x-lotus123|x-mdb|x-ms-excel|x-ms-project|x-mscardfile|x-msclip|x-msexcel|x-mspowerpoint|x-msproject|x-msword|x-msworks-db|x-msworks-wps|x-photoshop|x-postscript|x-powerpoint|x-ps|x-quark-express|x-rtf|x-vermeer-rpc|x-visio|x-vsd|x-wks|x-word|x-xls|x-xml|xhtml+xml|xls|xml)
	image/(photoshop|psd|x-photoshop|x-vsd)

Adding the profile to the virtual server
The final task is to add the profile to the iApp configuration. If you manually configured the system, simply attach the profile to the virtual server. If you used the iApp, use the following procedure.

To add the iRule to the virtual server

1. On the Main tab, expand **iApp** and then click **Application Services**.

2. Click the name of your HTTP Application service from the list.

3. On the Menu bar, click **Reconfigure**.

4. In the Delivery Optimization section, from the **Which compression profile do you want to use?** question, select the profile you just created.

5. Click the **Finished** button.

Next steps

After completing the iApp Template, the BIG-IP Application Services page opens for the HTTP service you just created. To see the list of all the configuration objects created to support the HTTP application, 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 HTTP 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 HTTP 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.

AVR statistics
If you have provisioned AVR, you can get application-level statistics for your HTTP Application Service.

To view AVR statistics

1. On the Main tab, expand **iApp** and then click **Application Services**.

2. From the Application Service List, click the HTTP service you just created.

3. On the Menu bar, click **Analytics**.

4. Use the tabs and the Menu bar to view different statistics for your iApp.

Object-level statistics
If you have not provisioned AVR, or want to view object-level statistics, use the following procedure.

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 version and had an existing Application Service that used the f5.http 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.http template has been significantly improved, and we strongly recommend you upgrade the source template to the new template available in v11.4 or 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, such as BIG-IP AFM in version 11.6.

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.http template.

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

Q: I configured the iApp template to use AFM to Reject or Log access connections from sources with low reputation scores, why isn't the system rejecting or logging those connections?

A: This is a known issue with the iApp template. If you are using BIG-IP AFM **_and_** IP Intelligence, and configured the iApp template to log or reject connection attempts from sources with a low reputation scores, you must configure the Blacklist categories manually before connections are logged or rejected. Use the following guidance.

1. Disable Strict Updates if you have not already done so:

    a. On the Main tab, expand **iApp** and then click **Application Services**.

    b. Click the name of your HTTP Application Service from the list.

    c. From the **Application Service** menu, select **Advanced**.

    d. In the **Strict Updates** row, clear the checkbox to disable Strict Updates.

    e. Click **Update**.

2. Click **Security > Network Firewall > IP Intelligence > Policies >** *(name-you-gave-the-iApp)*_**ip_intelligence**.

3. In the Blacklist Matching Policy area, from the **Blacklist Category** list, select a category that you want to log or reject, and then click **Add**.

4. Repeat step 3 to add all applicable blacklist categories.

5. Click **Update**.

6. You can optionally re-enable Strict Updates. Keep in mind, if you re-enter the iApp template and make changes to the configuration, you must perform this procedure again.

Appendix: Manual configuration table

We strongly recommend using the iApp template to configure the BIG-IP system for HTTP applications. Users familiar with the BIG-IP system can use the following table to manually configure the BIG-IP system. The table contains a list of configuration objects along with any non-default settings you should configure as a part of this deployment. Settings not mentioned in the table can be configured as applicable for your configuration. For specific instructions on configuring individual objects, see the Help tab or product manuals.

Health Monitors (*Local Traffic > Monitors*)	
Name	Type a unique name
Type	**HTTP** (**HTTPS** if you are deploying SSL Bridging)
Interval	**30** (recommended)
Timeout	**91** (recommended)
Send String	Optional. The iApp creates a send string based on HTTP version and FQDN: **GET /** HTTP/1.1\r\nHost: example.com\r\n
Receive String	Optional. Type the expected response from the Send String. If you leave this blank, the server is marked up if any response is received.
User Name[1]	If necessary, type a user name with access to your application. We recommend creating an account for use in this monitor. If you are using NTLM authentication and have a complex Active Directory environment where the username is not in the same domain as the Domain Controller being queried, you must prepend the domain name in uppercase letters (DOMAIN\) to the username in this field.
Password[1]	Type the associated password

Pools (*Local Traffic > Pools*)	
Name	Type a unique name
Health Monitor	Select the monitor you created above
Slow Ramp Time[2]	**300** (recommended)
Load Balancing Method	Choose a load balancing method. We recommend **Least Connections (Member)**
Address	Type the IP Address of the nodes
Service Port	Specify the Port; **80** is the most common HTTP port. Click **Add** to repeat Address and Service Port for all nodes.

Optional: AAM Application[3] (*Acceleration > Web Application*)	
Application Name	Type a unique name
Policy	**Generic Policy - Enhanced**
Requested Host	Type the Fully Qualified Domain Name (FQDN) of your application. Click **Add Host** to include additional hosts.

Profiles (*Local Traffic > Profiles*)		
HTTP (*Profiles > Services*)	Name	Type a unique name
	Parent Profile	**http**
	Rewrite Redirect[3]	**Matching**
TCP WAN (*Profiles > Protocol*)	Name	Type a unique name
	Parent Profile	**tcp-wan-optimized**
TCP LAN (*Profiles > Protocol*)	Name	Type a unique name
	Parent Profile	**tcp-lan-optimized**
Persistence (*Profiles > Persistence*)	Name	Type a unique name
	Persistence Type	**Cookie**
OneConnect (*Profiles > Other*)	Name	Type a unique name
	Parent Profile	**oneconnect**
Client SSL[4] (*Profiles > SSL*)	Name	Type a unique name
	Parent Profile	**clientssl**
	Certificate and Key	Select the Certificate and Key you imported from the associated list
Server SSL[5] (*Profiles > Other*)	Name	Type a unique name
	Parent Profile	**serverssl**
Web Acceleration (*Profiles > Services*)	Name	Type a unique name
	Parent Profile	**optimized-caching**
	WA Applications[2]	Enable the AAM Application you created

[1] User name and password are optional and only required if the health monitor should require credentials.
[2] You must select **Advanced** from the **Configuration** list for these options to appear
[3] *Optional. The BIG-IP AAM configuration is recommended, but optional.*
[4] *Only required if using the BIG-IP system for SSL Offload or SSL Bridging*
[5] *Only necessary if using the BIG-IP system for SSL Bridging or server-side encryption*

Profiles continued		
iSession[6] (Profiles > Services)	Name	Type a unique name
	Parent Profile	**isession**
HTTP Compression (Profiles > Services)	Name	Type a unique name
	Parent Profile	**wan-optimized-compression**
HTTP Compression (continued)	Content List--> Include List (Add each entry to the Content Type box and then click Include)	text/(css|html|javascript|json|plain|postscript|richtext|rtf|vnd\.wap\.wml|vnd\.wap\. wmlscript|wap|wml|x-component|x-vcalendar|x-vcard|xml)
		application/(css|css-stylesheet|doc|excel|javascript|json|lotus123|mdb|mpp|ms-excel|ms-powerpoint|ms-word|msaccess|msexcel|mspowerpoint|msproject|msword|photoshop|postscript|powerpoint|ps|psd|quarkexpress|rtf|txt|visio|vnd\.excel|vnd\.ms-access|vnd\.ms-excel|vnd\.ms-powerpoint|vnd\.ms-pps|vnd\.ms-project|vnd\.ms-word|vnd\.ms-works|vnd\.ms-works-db|vnd\.msaccess|vnd\.msexcel|vnd\.mspowerpoint|vnd\.msword|vnd\. powerpoint|vnd\.visio|vnd\.wap\.cmlscriptc|vnd\.wap\.wmlc|vnd\.wap\.xhtml\+xml|vnd\. word|vsd|winword|wks|word|x-excel|x-java-jnlp-file|x-javascript|x-json|x-lotus123|x-mdb|x-ms-excel|x-ms-project|x-mscardfile|x-msclip|x-msexcel|x-mspowerpoint|x-msproject|x-msword|x-msworks-db|x-msworks-wps|x-photoshop|x-postscript|x-powerpoint|x-ps|x-quark-express|x-rtf|x-vermeer-rpc|x-visio|x-vsd|x-wks|x-word|x-xls|x-xml|xhtml+xml|xls|xml)
		image/(photoshop|psd|x-photoshop|x-vsd)
Virtual Servers (*Local Traffic > Virtual Servers*)		
HTTP		
Name	Type a unique name.	
Address	Type the IP Address for the virtual server	
Service Port	**80**	
Protocol Profile (client)[1,2]	Select the WAN optimized TCP profile you created	
Protocol Profile (server)[1,2]	Select the LAN optimized TCP profile you created	
HTTP Profile[2]	Select the HTTP profile you created	
Web Acceleration profile[2]	Select the Web Acceleration profile you created	
HTTP Compression profile[2]	Select the HTTP Compression profile you created	
OneConnect[2]	Select the OneConnect profile you created	
Source Address Translation [3]	**Auto Map** (optional; see footnote [3])	
iSession profile[5]	If using BIG-IP AAM for symmetric optimization between systems, select the iSession profile you created.	
Default Pool[2]	Select the pool you created	
Persistence Profile[2]	Select the Persistence profile you created	
iRule[4]	If offloading SSL only: Enable the built-in **_sys_https_redirect irule**	
HTTPS[4]		
Name	Type a unique name.	
Address	Type the IP Address for the virtual server	
Service Port	**443**	
Protocol Profile (client)[1]	Select the WAN optimized TCP profile you created	
Protocol Profile (server)[1]	Select the LAN optimized TCP profile you created	
HTTP Profile	Select the HTTP profile you created above	
Web Acceleration profile	Select the Web Acceleration profile you created	
HTTP Compression profile	Select the HTTP Compression profile you created	
OneConnect	Select the OneConnect profile you created	
SSL Profile (Client)	Select the Client SSL profile you created	
SSL Profile (Server)[6]	If you created a Server SSL profile, select it from the list	
Source Address Translation[3]	**Auto Map** (optional; see footnote [3])	
iSession profile[5]	If using BIG-IP AAM for symmetric optimization between systems, select the iSession profile you created.	
Default Pool	Select the pool you created	
Persistence Profile	Select the Persistence profile you created	

[1] You must select **Advanced** from the **Configuration** list for these options to appear
[2] Do not enable these objects on the HTTP virtual server if offloading SSL. The HTTP virtual server is only used for redirecting users to the HTTPS virtual server.
[3] If using SNAT and expecting more than 64,000 simultaneous connections, you must configure a SNAT Pool with an IP address for each 64,000 simultaneous connections you expect. See the BIG-IP documentation on configuring SNAT Pools.
[4] Only necessary if offloading SSL or SSL Bridging
[5] Only necessary if using the BIG-IP AAM to provide symmetric optimization. Do not create/use this profile if you are deploying the BIG-IP on the server side of the WAN
[6] Only necessary if using the BIG-IP system for SSL Bridging or server-side encryption

Manually configuring the BIG-IP Advanced Firewall Module to secure your HTTP deployment

This section describes how to manually configure BIG-IP AFM, F5's Network Firewall module, to secure your HTTP deployment. BIG-IP AFM is particularly useful if you want to only allow access from specific clients or networks. Because this configuration can be complex, we recommend using the iApp template in version 11.6 and later to configure BIG-IP AFM.

Network Firewall settings

When configuring the BIG-IP Advanced Firewall Manager, you may want to configure your BIG-IP system to drop all traffic that you have not specifically allowed with firewall rules. This is known as *firewall mode*. By default, your BIG-IP system is set to default-accept, or *ADC mode*. Instructions for configuring your BIG-IP system, and the implications to consider, can be found on AskF5. For example, for BIG-IP v11.5: *http://support.f5.com/kb/en-us/products/big-ip-afm/manuals/product/network-firewall-policies-implementations-11-5-0/1.html*

If you have licensed IP Intelligence on the BIG-IP system, you can prohibit connections from sources with low reputation scores.

The following instructions cover a basic firewall configuration that is effective for the most common scenario of wanting to allow connections from a single trusted network. If you have complex requirements, such as the need to schedule different policies for different times of the day, or you want to create complicated rule or address lists, consult the BIG-IP AFM documentation. The basic steps for Policy and Rule creation apply to all scenarios.

To configure the BIG-IP AFM to allow connections from a single trusted network

1. Create a Network Firewall Policy:

    a. From the Configuration utility, click **Security > Network Firewall > Policies**, and then click **Create**.

    b. In the **Name** field, type a unique name for the policy, such as **HTTP-Policy**.

    c. Click **Finished**.

2. Create a rule to allow authorized hosts or networks to connect:

    a. Click **Security > Network Firewall > Policies**.

    b. Click the name of the policy you just created.

    c. In the Rule section (below the General Properties section), click the **Add** button.

    d. Leave the **Type** list set to Rule.

    e. From the **Order** list, select **First**. The Order list only appears in version 11.5 and later. In 11.4.x, you must reorder the rules from the Policy General Properties page.

    f. In the **Name** field, type a unique name, for instance **HTTP-traffic-Allowed**.

    g. Ensure the **State** list is set to **Enabled**.

    h. From the **Protocol** list, select **TCP**. Leave the box to the right of TCP set to **6**.

    i. In the **Source** section, from the **Address/Region** list, select **Specify**.
    You are now able to list the trusted source addresses for your connection.
    In the following example, we will configure a single subnet as trusted.

        • Select **Address**.

        • In the box, type the network address you want to allow, including netmask if more than a single host. Specify a network using CIDR notation, such as **10.0.0.0/24**.

        • Do not configure a source port.

        • Optional: If you want to limit inbound connections to a specific VLAN or Tunnel, from the **VLAN / Tunnel** list, select **Specify**, and then move the VLANs or tunnels that are allowed access to the Selected box.

        • Click **Add**.

        • Repeat these steps for additional hosts or networks. Use **Address List** or **Address Range** when appropriate.

    j. In the **Destination** section, leave the **Address/Region** and **Port** set to **Any**. Because you will be applying your policy to a virtual server that listens only on a single desired address and port, do not specify that information here.

k. If necessary, from the **Action** list, select **Accept**.

l. *Optional:* If you have configured a logging profile and want to log connections, from the **Logging** list, select **Enabled**. Typically, allowed connections do not need to be logged.

m. Click **Finished**.

3. Creating a firewall rule to block all other traffic
    The next task is to create a firewall rule to block all other traffic that you have not allowed. Although this is not a required step if your BIG-IP system is set to default deny (**Firewall mode**), it is required in default-accept (**ADC mode**), and is a good practice to always configure such a rule.

    a. Click **Security > Network Firewall > Policies**.

    b. Click the name of the policy you created in step 1.

    c. In the Rule section (below the General Properties section), click the **Add** button.

    d. Leave the **Type** list set to **Rule**.

    e. Leave the **Order** list, select **Last**.

    f. In the **Name** field, type a unique name, for example **HTTP-traffic-Prohibited**.

    g. Ensure the **State** list is set to **Enabled**.

    h. From the **Protocol** list, select **TCP**. Leave the box to the right of TCP set to **6**.

    i. In the **Source** section, leave all the lists set to **Any**

    j. From the **Action** list, select either **Drop** (to silently discard incoming connections) or **Reject** (to send a Destination Unreachable message to the sender).

    k. If you configured a logging profile as described in *Optional: Configuring the BIG-IP system to log network firewall events on page 32*, from the **Logging** list, select **Enabled**.

    l. Click **Finished**. You return to the Policy Properties page.

    m. On the Policy Properties page, in the Rules section, ensure the rule with the Action of Accept comes before the Drop or Reject rule you just created. If it does not, use the **Reorder** button and drag the rules into the correct order.

4. Apply Your Firewall Policy to your Virtual Server

    a. Click **Security > Network Firewall > Active Rules**.

    b. In the Rule section (below the General Properties section), click the **Add** button.

    c. From the **Context** list, select **Virtual Server**, and then select the virtual server you created for your HTTP traffic.

    d. From the **Type** list, select **Policy**, and then select the firewall policy you created.

    e. From the **Policy Type** list, select **Enforced**.

    f. Click **Finished**.

Optional: Assigning an IP Intelligence Policy to your HTTP virtual server

If you want to restrict access to your HTTP virtual server based on the reputation of the remote sender, you can enable and assign an IP Intelligence policy. This requires an IP intelligence license; contact your F5 Sales representative for more information.

It is outside the scope of this document to provide instructions on configuring an IP Intelligence Policy. Full documentation on enabling and configuring the IP Intelligence feature can be found on AskF5. For example, the manual for BIG-IP AFM v11.5 is: *https://support.f5.com/kb/en-us/products/big-ip-afm/manuals/product/network-firewall-policies-implementations-11-5-0/5.html*

After you have enabled and configured an IP Intelligence policy, use the following procedure to assign the policy to your HTTP virtual server:

To assign the IP intelligence policy to the HTTP virtual server

1. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.

2. Click the name of your HTTP virtual server.

3. From the **Security** menu, choose **Policies**.

4. Next to **IP Intelligence**, select **Enabled**, then select the IP intelligence policy to apply to traffic on the virtual server.

5. Click **Update**. The list screen and the updated item are displayed. The IP Intelligence policy is applied to traffic on the virtual server.

Optional: Configuring the BIG-IP system to log network firewall events

If you are using BIG-IP AFM, you have the option of logging network firewall events to one or more remote syslog servers (recommended) or to log events locally. You can either use an iApp template to create the logging profile, or create the logging profile manually.

For specific information on logging on the BIG-IP system, see the appropriate guide for your version. For example, for 11.5.0:

- Remote High-Speed Logging:
  https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos-implementations-11-5-0/22.html

- Local logging:
  https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos-concepts-11-5-0/11.html

Creating the logging profile using the iApp template

Use this section to create the logging profile using the logging profile iApp template. If you have not already downloaded the iApp template, see *https://devcentral.f5.com/wiki/iApp.F5-Remote-Logging-iApp.ashx*.

To configure the logging profile iApp

1. Log on to the BIG-IP system.

2. On the Main tab, click **iApp > Application Services**.

3. Click **Create**. The Template Selection page opens.

4. In the **Name** box, type a name. In our example, we use **logging-iapp_.**

5. From the **Template** list, select **f5.remote_logging.v<*latest-version*>**. The template opens

6. Use the following table for guidance on configuring the iApp template. Questions not mentioned in the table can be configured as applicable for your implementation.

Question	Your selection
Do you want to create a new pool of remote logging servers, or use an existing one?	Unless you have already created a pool on the BIG-IP system for your remote logging servers, select **Create a new pool**.
Which servers should be included in this pool?	Specify the IP addresses of your logging servers. Click **Add** to include more servers.
What port do the pool members use?	Specify the port used by your logging servers, typically **514**.
Do the pool members expect UDP or TCP connections?	**TCP**
Do you want to create a new monitor for this pool, or use an existing one?	Unless you have already created a health monitor for your pool of logging servers, select **Use a simple ICMP (ping) monitor**.
Do your log pool members require a specific log format?	If your logging servers require a specific format, select the appropriate format from the list.

7. Click **Finished**.

8. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.

9. Click the name of your HTTP virtual server.

10. From the **Security** menu, choose **Policies**.

11. Next to **Log Profile**, select **Enabled**, then select the Logging profile you created.

12. Click **Update**. The list screen and the updated item are displayed.

➡ **Note:** *The iApp template creates a log publisher and attaches it to the logging profile. If the publisher does not appear in the BIG-IP Configuration utility (GUI), you can verify the configuration by running the following command from the Traffic Management shell (tmsh):* `list security log profile <your profile name>`.

Creating logging profile manually

If you do not want to use the iApp template to create a logging profile, use this section for guidance on configuring the logging profile manually. You must have access to the tmsh command line to use this method.

To manually configure a logging profile

1. Use the following guidance for configuring a health monitor and load balancing pool for the logging servers.

BIG-IP LTM Object	Non-default settings/Notes	
Health Monitor (*Local Traffic -->Monitors*)	*Name*	Type a unique name
	Type	**ICMP**
	Interval	**30** (recommended)
	Timeout	**91** (recommended)
Pool (*Local Traffic -->Pools*)	*Name*	Type a unique name
	Health Monitor	Select the appropriate monitor you created
	Slow Ramp Time	**300**
	Load Balancing Method	Choose a load balancing method. We recommend **Least Connections (Member)**
	Address	Type the IP Address of a server.
	Service Port	Type the appropriate port, such as UDP port **514**, the port on which logging typically occurs. Click **Add**, and then repeat Address and Port for all nodes

2. Log into the BIG-IP system using the command line. Enter the tmsh shell, by typing **tmsh** from the prompt.

3. Create a Remote High Speed Log (HSL) destination:

   `(tmos)# create / sys log-config destination remote-high-speed-log [name] pool-name [specified pool] protocol [udp or tcp]`

4. If you have a specific log format requirement, create a format-specific log destination, and forward that to the previously-created HSL destination:

   `(tmos)# create / sys log-config destination [splunk|arcsight|remote-high-speed-log] [name] forward-to [HSL name]`

5. Create a log publisher:

   `(tmos)# create / sys log-config publisher [name] destinations add { [logdestination name] }`

6. Create the logging profile to tie everything together.
   If you chose to log allowed connections, include the green text (as in step 2 substep l in *To configure the BIG-IP AFM to allow connections from a single trusted network on page 30)*.
   If you set the rule to drop incoming connections, include the text in blue.
   If you chose to log IP intelligence events, include the text in red to add the parameter that sets the log publisher.

   `(tmos)# create / security log profile [name] network add { [name] { filter { log-acl-match-accept enabled log-acl-match-drop enabled log-acl-match-reject enabled } format { field-list { date_time action drop_reason protocol src_ip src_port dest_ip dest_port } type field-list } publisher [logpublisher name] } } ip-intelligence { log-publisher [logpublisher name] }`

Assigning the logging profile to the virtual server
The final task is to assign the logging profile to the virtual server.

To assign the logging profile to the HTTP virtual server

1. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.

2. Click the name of your HTTP virtual server.

3. From the **Security** menu, choose **Policies**.

4. Next to **Log Profile**, select **Enabled**, then select the Logging profile you created.

5. Click **Update**. The list screen and the updated item are displayed.

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 HTTP Application Service in the advanced configuration mode.

iSession
An iSession is an optimized connection between two BIG-IP systems.

iSession profile
An iSession profile defines the optimization parameters. WAN optimization requires an iSession profile, which specifies the optimization settings, such as compression and data deduplication. The iApp template uses the default isession profile.

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.
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.

local endpoint
The local endpoint is the BIG-IP system on which you are currently working. The systems must be set up symmetrically, so that a local endpoint connects to one or more remote endpoints.

network virtual server
A network virtual server is a virtual server whose IP address has no bits set in the host portion of the IP address (that is, the host portion of its IP address is 0, such as 192.168.1.0). This allows you to direct client traffic based on a range of destination IP addresses.

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, such as HTTP connections. 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 web 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. - Updated the iApp walk through section for additional options in 11.5, including the ability to have the BIG-IP system accept or deny client traffic from specific VLANs (see page 9), and SSL pass-through (see page 13).	01-31-2014
1.2	- Added the new section *Modifying the configuration produced by the iApp template if using 11.4.x or 11.5.x on page 24,* with a required change to the HTTP Compression profile. Added the same change to the manual configuration table. - Added support for BIG-IP v11.5.1.	06-03-2014
2.0	- Updated the guide for BIG-IP v11.6. This includes the ability to configure BIG-IP AFM using the iApp template. - Added the new section *Manually configuring the BIG-IP Advanced Firewall Module to secure your HTTP deployment on page 30*.	08-22-2014
2.1	Added the section *Troubleshooting on page 27* with a required modification to the configuration produced by the iApp template if using BIG-IP AFM and the IP Intelligence database to log or restrict traffic with low reputation scores.	11-04-2014
2.2	Added support for BIG-IP version 12.0	09-02-2015
2.3	Added BIG-IP ASM configuration information to the iApp walkthrough and the manual configuration table.	09-15-2015
2.4	Added missing "Advanced" icons to the relevant questions in the iApp walkthrough section.	01-05-2016
2.5	- Added a note to health monitor section in the iApp walkthough concerning when to prepend Domain information if using NTLM authentication. Added the same note to the manual configuration table. - Added support for BIG-IP version 12.1	05-18-2016
2.6	- Added a note in the ASM section of the iApp walkthrough stating that if using ASM in BIG-IP versions 12.1 and 12.1.1, you may need to create the ASM policy from the command line and not the BIG-IP GUI. iApp version 1.2.0rc4 contains the ability to create an ASM policy without having to manually configure the policy outside the iApp. - Added support for BIG-IP version 12.1.1	11-10-2016
2.7	Added support for BIG-IP version 13.0	02-22-2017
2.8	Updated this guide to include a reference on page 3 to the new, downloadable HTTP iApp template (v1.2.0) that is fully supported and available on downloads.f5.com.	05-12-2017