



Deploying the BIG-IP System for RADIUS Traffic Management

Welcome to the F5® deployment guide for RADIUS traffic management. This document provides guidance for configuring the BIG-IP® system version 11.4 and later for load balancing and intelligent traffic management for RADIUS implementations. BIG-IP version 11.0 introduced iApps™ Application templates, an extremely easy way to accurately configure the BIG-IP system for your RADIUS servers.

Products and Versions tested

Product	Version
BIG-IP LTM	11.4, 11.4.1, 11.5, 11.5.1, 11.6
RADIUS	Not applicable
RADIUS iApp template	System iApp that ships with v11.4 and later
Deployment Guide version	1.2 (see <i>Document Revision History on page 27</i>)

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

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

Contents

Why F5?	3
What is F5 iApp™?	3
Prerequisites and configuration notes	3
<hr/>	
Configuration scenarios	4
<hr/>	
Using this guide	4
<hr/>	
Preparing to use the iApp	5
<hr/>	
Configuring the BIG-IP iApp for RADIUS servers	6
Advanced options	6
Template Options	6
High Availability for Authentication and Authorization	7
Application Health for Authentication and Authorization	11
iRules for Authentication and Authorization	11
High Availability for Accounting	12
Application Health for Accounting	15
iRules for Accounting	15
Finished	16
<hr/>	
Optional: Modifying the iApp configuration if using MSCHAPv2	17
<hr/>	
Next steps	18
Modifying DNS settings to use the BIG-IP virtual server address	18
Upgrading an Application Service from previous version of the iApp template	19
<hr/>	
Appendix: Manual configuration table	20
<hr/>	
Appendix B: Test environment configuration information	21
<hr/>	
Glossary	24
<hr/>	
Document Revision History	27

Why F5?

The BIG-IP system provides a number of ways to accelerate, optimize, and scale RADIUS server deployments. The BIG-IP LTM uses an advanced health monitor that logs on to an RADIUS server to ensure traffic is only sent to available RADIUS servers.

What is F5 iApp™?

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

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

Prerequisites and configuration notes

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

- For this deployment guide, the BIG-IP system **must** be running version 11.4 or later. If you are using a previous version of the BIG-IP system, see the deployment guide index on F5.com. The configuration in this guide does not apply to previous versions.
- If you upgraded your BIG-IP system from a previous v11 version, and have an existing Application Service that used the f5.radius iApp template, see *Upgrading an Application Service from previous version of the iApp template on page 19*.
- This document provides guidance for using the iApp for RADIUS found in version 11.4 and later. For users familiar with the BIG-IP system, there is a manual configuration table at the end of this guide. However, we recommend using the iApp template.
- The BIG-IP health monitor created by the iApp requires an RADIUS user account. To check the health of the servers, the monitor uses this account to log in to RADIUS to verify server health. We recommend creating a new RADIUS user account for this health monitor.
- The RADIUS server must be configured to accept connections from BIG-IP Self IP address. Consult your RADIUS documentation for specific instructions.
In our example, we are using FreeRADIUS, so we add the BIG-IP address to the clients file, found in **/etc/freeradius/clients** with the following command syntax:

```
client 192.0.2.230 {  
    secret = testing123  
    shortname = bigip0  
}
```
- By default, the iApp configures Datagram load balancing. MSCHAPv2 (and other challenge/response authentication mechanisms) do not work with Datagram load balancing, due to multiple RADIUS packets per session. All packets in the conversation need to be delivered to the same server in order for this authentication mechanism to function correctly. If you are using MSCHAPv2 or another challenge/response authentication mechanism, see *Optional: Modifying the iApp configuration if using MSCHAPv2 and a BIG-IP version prior to 11.5 on page 17*.
- See *Appendix B: Test environment configuration information on page 21* and *Verifying successful RADIUS authN on page 22* for additional information.

Configuration scenarios

In this Deployment Guide, the BIG-IP system is optimally configured to optimize and direct traffic to RADIUS servers. This diagram shows a logical configuration example with a redundant pair of BIG-IP LTM devices in front of a group of RADIUS servers.




Figure 1: Logical configuration example

Using this guide

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

Using this guide to configure the App template

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

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

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

Manually configuring the BIG-IP system

Users already familiar with the BIG-IP system can use the manual configuration tables to configure the BIG-IP system for the RADIUS 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 20*.

Preparing to use the iApp

In order to use the iApp for RADIUS, 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 LTM Preparation table			
Basic/Advanced mode	In the iApp, you can configure your RADIUS implementation with F5 recommended settings (Basic mode) which are a result of extensive testing and tuning with RADIUS. Advanced mode gives you the to configure the BIG-IP system on a much more granular level, configuring specific options, or even using your own pre-built profiles or iRules. Basic and Advanced "configuration mode" is independent from the Basic/Advanced list at the very top of the template which only toggles the Device and Traffic Group options (see page 6)		
RADIUS Services	This iApp supports the following RADIUS services: Accounting, and Authentication and Authorization, or both.		
	Accounting		Authentication and Authorization
	<i>The is the address clients use to access the servers.</i>		<i>The is the address clients use to access the servers.</i>
	<i>IP address for the Virtual server:</i>		<i>IP address for the virtual server:</i>
	<i>Associated service port (default is 1646):</i>		<i>Associated service port (default is 1645):</i>
	<i>IP addresses of the RADIUS servers running the Accounting service</i>		<i>IP addresses of the RADIUS servers running the Auth/Auth service</i>
	1: 3: 5: 7: 9:	2: 4: 6: 8 10:	1: 2: 3: 4: 5: 6: 7: 8 9: 10:
Network	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 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
Application health	The iApp creates a health monitor that uses a RADIUS account to log into the server. We recommend a new user account specifically for use in the health monitor that is set to never expire. The monitor requires the following: <i>RADIUS user name:</i> <i>Password for that account:</i> <i>RADIUS Secret:</i> <i>Network Access Server (NAS) IP address:</i>		
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 RADIUS servers

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

Getting Started with the iApp for RADIUS servers

To begin the RADIUS iApp Template, use the following procedure.

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

Advanced options

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

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

Template Options

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

1. **Do you want to see inline help?**
Choose whether you want to see informational and help messages inline throughout the template, or if you would rather hide this inline help. If you are unsure, we recommend having the iApp display the inline help. Important and critical notes are always shown, no matter which selection you make.
 - ▶ **Yes, show inline help text**
Select this option to see all available inline help text.
 - ▶ **No, do not show inline help text**
If you are familiar with this iApp template, or with the BIG-IP system in general, select this option to hide the inline help text.
2. **Which configuration mode do you want to use?**
Select whether you want to use F5 recommended settings, or have more granular, advanced options presented.
 - ▶ **Basic - Use F5's recommended settings**
In basic configuration mode, options like load balancing method and parent profiles are all set automatically. The F5 recommended settings come as a result of extensive testing with web applications, so if you are unsure, choose Basic.
 - ▶ **Advanced - Configure advanced options**
In advanced configuration mode, you have more control over individual settings and objects, such as server-side optimizations and advanced options like Slow Ramp Time and Priority Group Activation. You can also choose to attach iRules you have previously created to the application service. The Advanced option provides more flexibility for experienced users.

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

3. **Which RADIUS services are you deploying at this time?**

Select which of the RADIUS services for which you are deploying this iApp. The iApp creates a virtual server for each of the services you select.

▶ **Accounting**

Choose this option if you are only deploying the RADIUS Accounting service at this time. The system creates a complete BIG-IP configuration for the Accounting service. Continue with [High Availability for Accounting on page 12](#).

▶ **Authentication and Authorization**

Choose this option if you are only deploying the RADIUS Authentication and Authorization service at this time. The system creates a complete BIG-IP configuration for the Authentication and Authorization service. Continue with High Availability for Authentication and Authorization, on this page.

▶ **Authentication and Authorization, and Accounting**

Choose this option if you want to deploy the iApp template for both Accounting and Authentication and Authorization. The system creates a virtual server for each service.

a. **Should the BIG-IP virtual servers use the same IP address?**

Select whether you want the system to use the same IP address for the BIG-IP virtual server for both the Accounting and Authentication and Authorization services.

▶ **No, the RADIUS services should have separate virtual IP addresses**

Select this option if you want to use unique IP addresses for each of the RADIUS services. You are asked for the appropriate IP address in the High Availability section for each service.

▶ **Yes, both RADIUS services should use the same virtual IP address**

Select this option if you want the system to use the same virtual IP address for both services. In this case, the BIG-IP system creates two virtual servers with this address, each listening on a different port.

i). **What IP address do you want to use for the virtual servers?**

Specify the IP address you want to use for the virtual servers. This IP address, combined with the port you specify in the High Availability section, become the BIG-IP virtual server address and port, which clients use to access the application. The system intercepts requests to the IP:Port combination and distributes them to the appropriate RADIUS servers.

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 servers.

ii). **What IP mask do you want applied to the virtual address?** **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 for the virtual server, you may leave this field blank.

High Availability for Authentication and Authorization

This section gathers information about your RADIUS Authentication and Authorization service, which will be used in the BIG-IP [virtual server](#) and [load balancing pool](#).

This section does not appear if you choose to only deploy the iApp for the Accounting only.

1. **What IP address do you want to use for the virtual server?**

This question does not appear if you choose both RADIUS services and the RADIUS services should use the same virtual IP address

Specify the IP address you want to use for the virtual server. This IP address, combined with the port you specify in #3 become the BIG-IP virtual server address and port, which clients use to access the application. The system intercepts requests to the IP:Port combination and distributes them to the appropriate RADIUS servers.

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 to the pool of servers.

2. **What IP mask do you want applied to the virtual address?** **Advanced**

This question only appears if you answered that the RADIUS services should have separate virtual IP addresses

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 for the virtual server, you may leave this field blank.

3. **What is the associated service port?**

Select the appropriate port(s) from the list. The system creates the virtual server using the port you specify. If you select both, the system creates two virtual servers, one on each port.

▶ **1645**

Select this option to enable the virtual server to listen on port 1645.

▶ **1812**

Select this option to enable the virtual server to listen on port 1812.

▶ **Both 1645 and 1812**

Select this option to enable the iApp to create two virtual servers for Authentication and Authorization, one that listens on port 1645 and the other on port 1812. Both of these virtual servers use the IP address you specified in #1.

4. **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 RADIUS servers, you can select it from the list. If you do select an existing pool, all of the rest of the questions in this section disappear.

▶ **Create a new pool**

Leave this default option to create a new load balancing pool and configure specific options.

a. **Are you using a challenge/response authentication method, such as MSCHAP?**

Select whether you are using an authentication method that uses challenges and responses, such as MSCHAPv2 or similar. This determines whether the BIG-IP system uses UDP datagram load balancing. Note that this is a setting on the UDP profile, and not the actual load balancing method.

▶ **Not using a challenge/response authentication**

Select this option (the default) if you are not using a challenge/response authentication method. The BIG-IP system enables Datagram load balancing on the UDP profile in this case.

▶ **Using a challenge/response authentication**

Select this option if you are using MSCHAP or other challenge/response authentication method. The system disables the Datagram load balancing setting on the UDP profile in this case.

b. **Which load balancing method do you want to use?** **Advanced**

Specify the load balancing method you want to use for this RADIUS pool. For RADIUS, we recommend the default, **Least Connections (member)**.

c. **Use a Slow Ramp time for newly added servers?** **Advanced**

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

Select whether you want to use a Slow Ramp time.

▶ **Use Slow Ramp**

Select this option for the system to implement Slow Ramp time for this pool.

i). *How many seconds should Slow Ramp time last?*

Specify a duration in seconds, for Slow Ramp. The time period you select for Slow Ramp is highly dependent on the speed of your server hardware and the behavior of your services. The default setting of 300 seconds (5 minutes) is very conservative in most cases.

▶ **Do not use Slow Ramp**

Select this option if you do not want to use Slow Ramp. If you select this option, we recommend you do not use the Least Connections load balancing method.

d. *Do you want give priority to specific groups of servers?* **Advanced**

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

▶ **Do not use Priority Group Activation**

Select this option if you do not want to enable Priority Group Activation.

▶ **Use Priority Group Activation**

Select this option if you want to enable Priority Group Activation.

You must add a priority to each server in the Priority box described in #e.

i). *What is the minimum number of active members in a group?*

Specify the minimum number of servers that must be active to continue sending traffic to the priority group. If the number of active servers falls below this minimum, traffic will be sent to the group of servers with the next highest priority group number.

e. *Which RADIUS servers are a part of this pool?*

Specify the IP address(es) of your RADIUS 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.

7. ***Where will the virtual servers be in relation to the RADIUS servers?***

Select whether your BIG-IP virtual servers are on the same subnet as your RADIUS servers, or on different subnets. This setting is used to determine the [SNAT](#) (secure NAT) and routing configuration.

▶ **BIG-IP virtual server IP and RADIUS servers are on the same subnet**

If the BIG-IP virtual servers and RADIUS 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 per server do you expect?*

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 per server**

Select this option if you expect fewer than 64,000 concurrent connections per 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 next section.

▶ **More than 64,000 concurrent connections per server**

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

i). 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.

1). What are the IP addresses 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 RADIUS 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 RADIUS servers?

If you chose different subnets, this question appears asking whether the 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.

- ▶ **RADIUS servers do not have a route to clients through the BIG-IP system**

If the 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.

i). How many connections per server do you expect?

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 per server**

Select this option if you expect fewer than 64,000 concurrent connections per 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 next section.

- **More than 64,000 concurrent connections per server**

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.

1). 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.

Application Health for Authentication and Authorization

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

1. **Create a new health monitor or use an existing one?**

Application health monitors are used to verify the Authentication and Authorization service of the RADIUS servers is available and functioning.

Unless you have requirements for configuring other options not in the following list of questions, we recommend allowing the iApp to create a new RADIUS 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 a health monitor, select it from the list. Continue with the next section.

▶ **Create a new health monitor**

If you want the iApp to create a new RADIUS monitor, continue with the following.

a. **How many seconds should pass between health checks?**

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

b. **What user account should this health monitor use to log into the RADIUS servers?**

The health monitor requires an RADIUS user account, which the BIG-IP uses to log on to the server as a part of the health check. We recommend creating a new account specifically for this monitor.

c. **What is the password for the specified user account?**

Type the associated password.

d. **What is the secret?**

Specify the RADIUS secret the monitor will need to access the RADIUS server.

e. **What is the NAS IP address?**

Specify the IP address of the Network Address Server (NAS) that is a part of your RADIUS implementation.

iRules for Authentication and Authorization

In this section, you can add custom iRules to the Authentication and Authorization service of the RADIUS 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 the Authentication and Authorization service of the RADIUS implementation.

**Warning**

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

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 RADIUS 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.

High Availability for Accounting

This section gathers information about your RADIUS Accounting service, which is used in the BIG-IP [virtual server](#) and [load balancing pool](#).

This section does not appear if you choose to only deploy the iApp for the Authentication and Authorization only.

1. **What IP address do you want to use for the virtual server?**

This question does not appear if you choose both RADIUS services and the RADIUS services should use the same virtual IP address

Specify the IP address you want to use for the virtual server. This IP address, combined with the port you specify in #3 become the BIG-IP virtual server address and port, which clients use to access the application. The system intercepts requests to the IP:Port combination and distributes them to the appropriate RADIUS servers.

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 servers.

2. **What IP mask do you want applied to the virtual address?** **Advanced**

This question only appears if you answered that the RADIUS services should have separate virtual IP addresses

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 for the virtual server, you may leave this field blank.

3. **What is the associated service port?**

Select the appropriate port(s) from the list. The system creates the virtual server using the port you specify. If you select both, the system creates two virtual servers, one on each port.

▶ **1646**

Select this option to enable the virtual server to listen on port 1646.

▶ **1813**

Select this option to enable the virtual server to listen on port 1813.

▶ **Both 1646 and 1813**

Select this option to enable the iApp to create two virtual servers for Accounting, one that listens on port 1646 and the other on port 1816. Both of these virtual servers use the IP address you specified in #1.

4. **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 RADIUS servers, you can select it from the list.

If you do select an existing pool, all of the rest of the questions in this section disappear.

▶ **Create a new pool**

Leave this default option to create a new load balancing pool and configure specific options.

a. **Which load balancing method do you want to use?** **Advanced**

Specify the load balancing method you want to use for this RADIUS pool. For RADIUS, we recommend the default, **Least Connections (member)**.

b. Use a Slow Ramp time for newly added servers? **Advanced**

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

Select whether you want to use a Slow Ramp time.

▶ **Use Slow Ramp**

Select this option for the system to implement Slow Ramp time for this pool.

i). How many seconds should Slow Ramp time last?

Specify a duration in seconds, for Slow Ramp. The time period you select for Slow Ramp is highly dependent on the speed of your server hardware and the behavior of your services. The default setting of 300 seconds (5 minutes) is very conservative in most cases.

▶ **Do not use Slow Ramp**

Select this option if you do not want to use Slow Ramp. If you select this option, we recommend you do not use the Least Connections load balancing method.

c. Do you want give priority to specific groups of servers? **Advanced**

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

▶ **Do not use Priority Group Activation**

Select this option if you do not want to enable Priority Group Activation.

▶ **Use Priority Group Activation**

Select this option if you want to enable Priority Group Activation.

You must add a priority to each server in the Priority box described in #e.

i). What is the minimum number of active members in a group?

Specify the minimum number of servers that must be active to continue sending traffic to the priority group. If the number of active servers falls below this minimum, traffic will be sent to the group of servers with the next highest priority group number.

d. Which RADIUS servers are a part of this pool?

Specify the IP address(es) of your RADIUS 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.

7. Where will the virtual servers be in relation to the RADIUS servers?

Select whether your BIG-IP virtual servers are on the same subnet as your RADIUS servers, or on different subnets. This setting is used to determine the [SNAT](#) (secure NAT) and routing configuration.

▶ **BIG-IP virtual server IP and RADIUS servers are on the same subnet**

If the BIG-IP virtual servers and RADIUS 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 per server do you expect?

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 per server**

Select this option if you expect fewer than 64,000 concurrent connections per 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 next section.

▶ **More than 64,000 concurrent connections per server**

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

i). 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.

1). What are the IP addresses 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 RADIUS 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 RADIUS servers?

If you chose different subnets, this question appears asking whether the 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.

▶ **RADIUS servers do not have a route to clients through the BIG-IP system**

If the 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.

i). How many connections per server do you expect?

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 per server**

Select this option if you expect fewer than 64,000 concurrent connections per 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 next section.

- **More than 64,000 concurrent connections per server**

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.

1). *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.

Application Health for Accounting

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

1. ***Create a new health monitor or use an existing one?***

Application health monitors are used to verify the Accounting service of the RADIUS servers is available and functioning.

Unless you have requirements for configuring other options not in the following list of questions, we recommend allowing the iApp to create a new RADIUS 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 a health monitor, select it from the list. Continue with the next section.

▶ **Create a new health monitor**

If you want the iApp to create a new RADIUS monitor, continue with the following.

a. *How many seconds should pass between health checks?*

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

b. *What user account should this health monitor use to log into the RADIUS servers?*

The health monitor requires an RADIUS user account, which the BIG-IP uses to log on to the server as a part of the health check. We recommend creating a new account specifically for this monitor.

c. *What is the password for the specified user account?*

Type the associated password.

d. *What is the secret?*

Specify the RADIUS secret the monitor will need to access the RADIUS server.

e. *What is the NAS IP address?*

Specify the IP address of the Network Address Server (NAS) that is a part of your RADIUS implementation.

iRules for Accounting

In this section, you can add custom iRules to the Accounting service of the RADIUS 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 you have preexisting iRules you want to add to the Accounting service of the RADIUS implementation.



Warning

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

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 RADIUS servers, from the **Options** box, click the name of the applicable iRule(s) and then click the Add (<<) button to move them to the **Selected** box.

Finished

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

Optional: Modifying the iApp configuration if using MSCHAPv2 and a BIG-IP version prior to 11.5

If you are using MSCHAPv2 or another challenge/response authentication mechanism, you must disable Datagram load balancing (currently configured by the iApp in versions prior to 11.5 by default). To modify the Datagram load balancing setting, you must first disable the Strict Updates feature on the iApp.

Disabling the Strict Updates feature

Before modifying the configuration produced by the iApp, you must turn off the Strict Updates feature. By turning off Strict Updates, if you re-enter the iApp template and modify the configuration within the iApp, you will have to make this change again manually.

To turn off Strict Updates

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. Click the name of your RADIUS Application service from the list.
3. From the **Application Service** list, select **Advanced**.
4. In the **Strict Updates** row, clear the check from the box to disable Strict Updates.
5. Click the **Update** button.

Modifying the Datagram load balancing setting

The next task is to modify the UDP profile created by the iApp to disable Datagram load balancing.

To modify the UDP profile

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. Click the name of your RADIUS Application service from the list.
3. On the Menu bar, click **Components**. The BIG-IP objects for the iApp appear.
4. From the list, click the name of the UDP profile that was created by the iApp. This profile is preceded by the name you gave the iApp, followed by **accounting_udp** or **auth_auth_udp**. If you selected both RADIUS services, both are present; click one of the two.
5. Click to clear check from the **Datagram LB Enabled** box.
6. Click the **Update** button.
7. If necessary, repeat for the other RADIUS service.
8. *Optional:* After modifying the template, we recommend turning Strict Updates back on. However, if you modify the iApp in the future, you must modify the UDP profile again.
To turn on Strict Updates, use the procedure above for turning off Strict Updates, but in Step 4, click a check in the box to Enable Strict Updates.

Next steps

After completing the iApp Template, the BIG-IP Application Services page opens for the RADIUS service you just created. To see the list of all the configuration objects created to support RADIUS, 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 RADIUS 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 RADIUS Application Service from the list.
3. On the Menu bar, click **Reconfigure**.
4. Make the necessary modifications to the template.
5. Click the **Finished** button.

Viewing statistics

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

Object-level statistics

Use the following procedure to view statistics.

To view object-level statistics

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

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

Upgrading an Application Service from previous version of the iApp template

If you upgraded your BIG-IP system from a previous v11 version and had an existing Application Service that used the f5.radius 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.radius template has been significantly improved, and we strongly recommend you upgrade the source template to the new template available in v11.4 and later.

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

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

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



Warning

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

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

Appendix: Manual configuration table

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

BIG-IP LTM Object	Non-default settings/Notes	
Health Monitor (Local Traffic-->Monitors)	Authentication and Authorization	
	Name	Type a unique name
	Type	RADIUS
	Interval	30 (recommended)
	Timeout	91 (recommended)
	User Name	Specify the RADIUS user name you want to use for this monitor. We suggest creating a new account specifically for this health check with no other permissions and set to never expire.
	Password	Type the associated password
	Secret	Type the RADIUS secret for your implementation
	NAS IP Address	Type the IP address of the Network Access Server in your RADIUS implementation.
	Accounting	
	Name	Type a unique name
	Type	RADIUS Accounting
	Interval	30 (recommended)
	Timeout	91 (recommended)
User Name	Specify the RADIUS user name you want to use for this monitor. We suggest creating a new account specifically for this health check with no other permissions and set to never expire.	
Secret	Type the RADIUS secret for your implementation	
NAS IP Address	Type the IP address of the Network Access Server in your RADIUS implementation.	
Pool (Local Traffic -->Pools)	Name	Type a unique name
	Health Monitor	Select the appropriate monitor you created above
	Slow Ramp Time¹	300
	Load Balancing Method	Choose a load balancing method. We recommend Least Connections (member)
	Address	Type the IP Address of the RADIUS nodes
	Service Port	Specify the appropriate port (1646 or 1813 for Accounting and 1645 or 1812 Authentication and Authorization). Click Add to repeat Address and Service Port for all nodes.
	If you are deploying both RADIUS services, repeat this section to create a pool for the other service.	
Profiles (Local Traffic-->Profiles)	Name	Type a unique name
	UDP (Profiles > Protocol)	Parent Profile UDP
		Datagram LB Enabled (If using MSCHAPv2 or similar challenge/response auth method, leave Datagram LB Disabled).
Virtual Server (Local Traffic --> Virtual Servers)	Name	Type a unique name.
	Address	Type the IP Address for the virtual server
	Service Port	1646 or 1813 for Accounting and 1645 or 1812 Authentication and Authorization
	Protocol Profile (Client)¹	Select the UDP profile you created above
	Source Address Translation ⁴	Auto Map (optional; see footnote ²)
	Default Pool	Select the pool you created above
	If you are deploying both RADIUS services, repeat this section to create a virtual server for the other service. You may also need to create a virtual server for both ports for each service (for example, two Accounting virtual servers using the same IP address, but one on port 1646 and one on port 1813), so you could have a total of four virtual servers, depending on your configuration.	

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

² If want to use SNAT, and you have a large deployment 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.

Appendix B: Test environment configuration information

The following information shows the configuration of the non-F5 devices used in our test environment.

Participating Nodes

- FreeRADIUS server listening on: 192.168.12.29/24:1812
- BIG-IP LTM
 - Self IP VLAN 12: 192.168.12.230/24
 - Self IP VLAN 245: 192.168.245.230/24
- Dell 5424 switch
 - VLAN 245 IP 192.168.245.201/24
 - ethernet g21 configured for 802.1X authentication
 - configured to authenticate 802.1X via RADIUS server @ 192.168.245.129:1812
- Linux supplicant host
 - Attached directly to switchport g21 via eth3
 - using wpa_supplicant to issue 802.1X EAP frames to switchport

Configure FreeRADIUS (Debian Squeeze)

- Add BIG-IP to clients file (/etc/freeradius/clients.conf):

```
client 192.168.12.230 {
    secret = testing123
    shortname = bigip0
}
```
- Add credentials to users file (/etc/freeradius/users)

```
proliant0eth3 Cleartext-Password := "testing"
steve Cleartext-Password := "testing" # exists in default config
```
- Start freeradius in debug mode
 - \$ sudo /etc/init.d/freeradius stop
 - \$ sudo /usr/sbin/freeradius -X

Configure Dell 5424 Switch

- dot1x system-auth-control

```
interface ethernet g21
    dot1x port-control auto
    dot1x re-authentication
    dot1x max-req 10
    dot1x timeout re-authperiod 300
    dot1x timeout quiet-period 1
exit
```
- interface vlan 245

```
ip address 192.168.245.201 255.255.255.0
exit
```
- radius-server host 192.168.245.129 auth-port 1812 key testing123 usage dot1x
- aaa authentication dot1x default radius

Configure 802.1X supplicant (Debian Squeeze)

- create wpa_supplicant-eth3 (man wpa_supplicant.conf)

```
ctrl_interface=/var/run/wpa_supplicant
ap_scan=0
network={
    key_mgmt=IEEE8021X
    eap=MD5
    identity="proliant0eth3"
    password="testing"
    eapol_flags=0
}
```

- Start wpa_supplicant in debug mode on interface attached to g21

```
$ sudo apt-get -y install wpasupplicant
$ sudo wpa_supplicant -d -Dwired -ieth3 -c/home/cjac/tmp/wpa_supplicant-eth3
```

Verifying successful RADIUS authN

The following is the output to the FreeRADIUS debug console during a successful EAP/MD5 authentication request in our example

```
rad_recv: Access-Request packet from host 192.168.12.230 port 49158, id=0, length=91
    NAS-IP-Address = 192.168.245.201
    NAS-Port-Type = Ethernet
    NAS-Port = 21
    User-Name = "proliant0eth3"
    EAP-Message = 0x02b300120170726f6c69616e743065746833
    Message-Authenticator = 0xf2a175759c8f1c09847530924206f050
# Executing section authorize from file /etc/freeradius/sites-enabled/default
+- entering group authorize {...}
++[preprocess] returns ok
++[chap] returns noop
++[mschap] returns noop
++[digest] returns noop
[suffix] No '@' in User-Name = "proliant0eth3", looking up realm NULL
[suffix] No such realm "NULL"
++[suffix] returns noop
[eap] EAP packet type response id 179 length 18
[eap] No EAP Start, assuming it's an on-going EAP conversation
++[eap] returns updated
[files] users: Matched entry proliant0eth3 at line 90
++[files] returns ok
++[expiration] returns noop
++[logintime] returns noop
[pap] WARNING: Auth-Type already set. Not setting to PAP
++[pap] returns noop
Found Auth-Type = EAP
# Executing group from file /etc/freeradius/sites-enabled/default
+- entering group authenticate {...}
[eap] EAP Identity
[eap] processing type md5
rlm_eap_md5: Issuing Challenge
++[eap] returns handled
```

```
Sending Access-Challenge of id 0 to 192.168.12.230 port 49158
    EAP-Message = 0x01b400160410f4e0d93077f6bee67f014150792e4312
    Message-Authenticator = 0x00000000000000000000000000000000
    State = 0x233c09a423880df339ce99821b06cde4
Finished request 116.
Going to the next request
Waking up in 4.9 seconds.
rad_recv: Access-Request packet from host 192.168.12.230 port 49158, id=0, length=113
# Executing section authorize from file /etc/freeradius/sites-enabled/default
+- entering group authorize {...}
++[preprocess] returns ok
++[chap] returns noop
++[mschap] returns noop
++[digest] returns noop
[suffix] No '@' in User-Name = "proliant0eth3", looking up realm NULL
[suffix] No such realm "NULL"
++[suffix] returns noop
[eap] EAP packet type response id 180 length 22
[eap] No EAP Start, assuming it's an on-going EAP conversation
++[eap] returns updated
[files] users: Matched entry proliant0eth3 at line 90
++[files] returns ok
++[expiration] returns noop
++[logintime] returns noop
[pap] WARNING: Auth-Type already set. Not setting to PAP
++[pap] returns noop
Found Auth-Type = EAP
# Executing group from file /etc/freeradius/sites-enabled/default
+- entering group authenticate {...}
[eap] Request found, released from the list
[eap] EAP/md5
[eap] processing type md5
[eap] Freeing handler
++[eap] returns ok
# Executing section post-auth from file /etc/freeradius/sites-enabled/default
+- entering group post-auth {...}
++[exec] returns noop
Sending Access-Accept of id 0 to 192.168.12.230 port 49158
    EAP-Message = 0x03b40004
    Message-Authenticator = 0x00000000000000000000000000000000
    User-Name = "proliant0eth3"
Finished request 117.
```

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 RADIUS application service in the advanced configuration mode.

load balancing method

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

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

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

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

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

load balancing pool

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

profile

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

self IP address

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

SNAT

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

SNAT pool

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

virtual server

A virtual server is a traffic-management object on the BIG-IP system that is represented by an IP address and a service port. This is the address clients use to connect to the RADIUS 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	<ul style="list-style-type: none"> - Added support for BIG-IP v11.4.1 and 11.5. - The updated iApp template in 11.5 contains the following fixes and additions: <ul style="list-style-type: none"> * The iApp no longer displays an error when configuring a SNAT Pool. * Added a question asking if a challenge/response authentication method such as MSCHAP is being used (see page 8). - Modified the section <i>Optional: Modifying the iApp configuration if using MSCHAPv2 and a BIG-IP version prior to 11.5 on page 17</i> to apply to versions prior to 11.5 only. 	01-31-2014
1.2	Added support for BIG-IP v11.5.1 and 11.6.	08-25-2014

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

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

F5 Networks
Asia-Pacific
apacinfo@f5.com

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

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

