BIG-IP® Local Traffic Controller Installation Guide

version 3.3

Service and Support Information

Product Version

This manual applies to version 3.3 of the BIG-IP® Local Traffic Controller.

Obtaining Technical Support

Web	tech.f5.com
Phone	(206) 272-6888
Fax	(206) 272-6802
Email (support issues)	support@f5.com
Email (suggestions)	feedback@f5.com

Contacting F5 Networks

Web	www.f5.com
Toll-free phone	(888) 961-7242
Corporate phone	(206) 272-5555
Fax	(206) 272-5556
Email	sales@f5.com
Mailing Address	501 Elliott Avenue West Seattle, Washington 98119

Legal Notices

Copyright

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described herein. F5 reserves the right to change specifications at any time without notice.

Trademarks

F5, BIG-IP, and 3-DNS are registered trademarks of F5 Networks, Inc. SEE-IT, GLOBAL-SITE, EDGE-FX, and FireGuard are trademarks of F5 Networks, Inc. Other product and company names are registered trademarks or trademarks of their respective holders.

Export Regulation Notice

The BIG-IP® Local Traffic Controller may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this BIG-IP® Local Traffic Controller from the United States.

Export Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

FCC Compliance

This equipment generates, uses, and may emit radio frequency energy. The equipment has been type tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules, which are designed to provide reasonable protection against such radio frequency interference.

Operation of this equipment in a residential area may cause interference, in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This class A digital apparatus complies with Canadian I CES-003.

Standards Compliance

The product conforms to ANSI/UL Std 1950 and Certified to CAN/CSA Std. C22.2 No. 950.

Acknowledgments

This product includes software developed by the University of California, Berkeley and its contributors.

This product includes software developed by the Computer Systems Engineering Group at the Lawrence Berkeley Laboratory.

This product includes software developed by the NetBSD Foundation, Inc. and its contributors.

This product includes software developed by Christopher G. Demetriou for the NetBSD Project.

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

This product includes software developed by Jochen Pohl for the NetBSD Project.

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

This product includes software developed by SigmaSoft, Th. Lockert.

This product includes software developed for the NetBSD Project by Jason R. Thorpe.

This product includes software developed by Jason R. Thorpe for And Communications, http://www.and.com.

This product includes software developed for the NetBSD Project by Frank Van der Linden.

This product includes software developed for the NetBSD Project by John M. Vinopal.

This product includes software developed by Christos Zoulas.

This product includes software developed by Charles Hannum.

This product includes software developed by Charles Hannum, by the University of Vermont and State Agricultural College and Garrett A. Wollman, by William F. Jolitz, and by the University of California, Berkeley, Lawrence Berkeley Laboratory, and its contributors.

This product includes software developed by the University of Vermont and State Agricultural College and Garrett A. Wollman.

In the following statement, "This software" refers to the Mitsumi CD-ROM driver: This software was developed by Holger Veit and Brian Moore for use with "386BSD" and similar operating systems. "Similar operating systems" includes mainly non-profit oriented systems for research and education, including but not restricted to "NetBSD," "FreeBSD," "Mach" (by CMU).

In the following statement, "This software" refers to the parallel port driver: This software is a component of "386BSD" developed by William F. Jolitz, TeleMuse.

This product includes software developed by the Apache Group for use in the Apache HTTP server project

(http://www.apache.org/).

This product includes software developed by Darren Reed. (© 1993-1998 by Darren Reed).

This product includes software licensed from Richard H. Porter under the GNU Library General Public License (© 1998, Red Hat Software), www.gnu.org/copyleft/lgpl.html.

This product includes the standard version of Perl software licensed under the Perl Artistic License (© 1997, 1998 Tom Christiansen and Nathan Torkington). All rights reserved. You may find the most current standard version of Perl at http://www.perl.com.

F5 Networks Limited Warranty

This warranty will apply to any sale of goods or services or license of software (collectively, "Products") from F5 Networks, Inc. ("F5"). Any additional or different terms including terms in any purchase order or order confirmation will have no effect unless expressly agreed to in writing by F5. Any software provided to a Customer is subject to the terms of the End User License Agreement delivered with the Product.

Limited Warranty

Software. F5 warrants that for a period of 90 days from the date of shipment: (a) the media on which the software is furnished will be free of defects in materials and workmanship under normal use; and (b) the software substantially conforms to its published specifications. Except for the foregoing, the software is provided AS IS.

In no event does F5 warrant that the Software is error free, that the Product will operate with any software or hardware other than that provided by F5 or specified in the documentation, or that the Product will satisfy Purchaser's own specific requirements.

Hardware. F5 warrants that the hardware component of any Product will, for a period of one year from the date of shipment from F5, be free from defects in material and workmanship under normal use.

Remedy. Purchaser's exclusive remedy and the entire liability of F5 under this limited warranty and any other guarantee made by F5 is, at F5's option, to repair or replace any Product or component that fails during the warranty period at no cost to Purchaser. Products returned to F5 must be pre-authorized by F5 with a Return Material Authorization (RMA) number marked on the outside of the package, and sent prepaid, insured, and packaged appropriately for safe shipment. The repaired or replaced item will be shipped to Purchaser, at F5's expense, no later than 7 days after receipt by F5. Title to any returned Products or components will transfer to F5 upon receipt. F5 will replace defective media or documentation or, at its option, undertake reasonable efforts to modify the software to correct any substantial non-conformance with the specifications.

Restrictions. The foregoing limited warranties extend only to the original Purchaser, and do not apply if a Product (a) has been altered, except by F5, (b) has not been installed, operated, repaired, or maintained in accordance with F5's instructions, (c) has been subjected to abnormal physical or electrical stress, misuse, negligence, or accident, or (d) has been operated outside of the environmental specifications for the Product. F5's limited software warranty does not apply to software corrections or upgrades.

Support, Upgrades. F5 provides software telephone support services at no charge for 90 days following the installation of any Product: Monday through Friday, from 6 a.m. to 6 p.m. Pacific time, excluding F5's holidays. Such support will consist of responding to trouble calls as reasonably required to make the Product perform as described in the Specifications. For advisory help requests, which are calls of a more consultative nature than a standard trouble call, F5 will provide up to two hours of telephone service at no charge. Additional service for advisory help requests may be purchased at F5 Networks' then-current standard service fee. During this initial 90

day period, Customer is entitled, at no charge, to updated versions of covered software such as bug fixes, and incremental enhancements as designated by minor revision increases. In addition, Customer will receive special pricing on upgraded versions of covered Products such as new clients, new modules, and major enhancements designated by major revision increases. Customer may purchase a Maintenance Agreement for enhanced maintenance and support services.

DISCLAIMER; LIMITATION OF REMEDY: EXCEPT FOR THE WARRANTIES SPECIFICALLY DESCRIBED HEREIN, F5 DOES NOT MAKE ANY GUARANTEE OR WARRANTY, EXPRESS OR IMPLIED, WITH RESPECT TO PRODUCTS, SPECIFICATIONS, SUPPORT, SERVICE, OR ANYTHING ELSE. F5 HAS NOT AUTHORIZED ANYONE TO MAKE ANY REPRESENTATION OR WARRANTY OTHER THAN AS PROVIDED ABOVE. F5 DISCLAIMS ANY AND ALL WARRANTIES AND GUARANTEES, EXPRESS, IMPLIED, OR OTHERWISE, ARISING WITH RESPECT TO THE PRODUCTS OR SERVICES DELIVERED HEREUNDER, INCLUDING BUT NOT LIMITED TO THE WARRANTY OF MERCHANTABILITY, THE WARRANTY OF FITNESS FOR A PARTICULAR PURPOSE, AND ANY WARRANTY OF NON-INFRINGEMENT OF THE INTELLECTUAL PROPERTY RIGHTS OF ANY THIRD PARTY. F5 WILL HAVE NO OBLIGATION OR LIABILITY, WHETHER ARISING IN CONTRACT (INCLUDING WARRANTY), TORT (INCLUDING ACTIVE, PASSIVE, OR IMPUTED NEGLIGENCE, STRICT LIABILITY, OR PRODUCT LIABILITY), OR OTHERWISE FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR INDIRECT DAMAGES INCLUDING BUT NOT LIMITED TO LOSS OF USE, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF REVENUE, LOSS OF BUSINESS, OR OTHER FINANCIAL LOSS ARISING OUT OF OR IN CONNECTION WITH ANY OF THE PRODUCTS OR OTHER GOODS OR SERVICES FURNISHED TO CUSTOMER BY F5, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

End-user Software License

IMPORTANT - READ BEFORE INSTALLING OR OPERATING THIS PRODUCT

CAREFULLY READ THE TERMS AND CONDITIONS OF THIS LICENSE BEFORE INSTALLING OR OPERATING THIS PRODUCT - BY INSTALLING, OPERATING OR KEEPING THIS PRODUCT FOR MORE THAN THIRTY DAYS AFTER DELIVERY YOU INDICATE YOUR ACCEPTANCE OF THESE TERMS AND CONDITIONS. IF YOU DO NOT AGREE WITH THE TERMS AND CONDITIONS OF THIS AGREEMENT, PROMPTLY CONTACT F5 NETWORKS, INC. ("F5") TO ARRANGE FOR RETURN OF THE PRODUCT FOR A REFUND.

1. Scope. This License applies to the software component ("Software") of the F5 product identified above ("Product") and any corrections, updates, new releases and new versions of such software. This License is a legal agreement between F5 and the single entity ("Licensee") that has acquired Software from F5 under applicable terms and conditions.

2. License Grant. Subject to the terms of this License, F5 grants to Licensee a non-exclusive, non-transferable license to use the Software in object code form with an unlimited number of servers. Other than as specifically described herein, no right or license is granted to Licensee to any of F5's trademarks, copyrights, or other intellectual property rights. The Software incorporates certain third party software, which is used subject to licenses from the respective owners. The protections given to F5 under this License also apply to the suppliers of this third party software, who are intended third party beneficiaries of this License.

3. Restrictions. The Software, documentation, and the associated copyrights and other intellectual

property rights are owned by F5 or its licensors, and are protected by law and international treaties. Licensee may not copy or reproduce the Software, and may not copy or translate the written materials without F5's prior, written consent. Licensee may not copy, modify, reverse compile, or reverse engineer the Software, or sell, sub-license, rent, or transfer the Software or any associated documentation to any third party.

4. Export Control. F5's standard Software incorporates cryptographic software. Licensee agrees to comply with the Export Administration Act, the Export Control Act, all regulations promulgated under such Acts, and all other US government regulations relating to the export of technical data and equipment and products produced therefrom, which are applicable to Licensee. In countries other than the US, Licensee agrees to comply with the local regulations regarding importing, exporting, or using cryptographic software.

5. Limited Warranty. F5 warrants that for a period of 90 days from the date of shipment: (i) the media on which the Software is furnished will be free of defects in materials and workmanship under normal use; and (ii) the Software substantially conforms to its published specifications. Except for the foregoing, the Software is provided AS IS. In no event does F5 warrant that the Software is error free, that it will operate with any software or hardware other than that provided by F5 or specified in the documentation, or that the Software will satisfy Licensee's own specific requirements.

   a. Remedy. Licensee's exclusive remedy and the entire liability of F5 under this limited warranty and any other guarantee made by F5 is, at F5's option, to repair or replace any Software that fails during the warranty period at no cost to Licensee. Any Product returned to F5 must be pre-authorized by F5 with a Return Material Authorization (RMA) number marked on the outside of the package, and sent prepaid, insured and packaged appropriately for safe shipment. The repaired or replaced item will be shipped to Licensee, at F5's expense, no later than 7 days after receipt by F5. Title to any returned Products or components will transfer to F5 upon receipt. F5 will replace defective media or documentation or, at its option, undertake reasonable efforts to modify the Software to correct any substantial non-conformance with the specifications.

   b. Restrictions. The foregoing limited warranties extend only to the original Licensee, and do not apply if the Software or the Product (a) has been altered, except by F5, (b) has not been installed, operated, repaired, or maintained in accordance with F5's instructions, (c) has been subjected to abnormal physical or electrical stress, misuse, negligence or accident or (d) has been operated outside of the environmental specifications for the Product. F5's limited software warranty does not apply to software corrections or upgrades.

6. Infringement Indemnity. F5 will, at its expense, defend any suit brought against Licensee based upon a claim that the Software as delivered by F5 directly infringes a valid patent or copyright. F5 will pay costs and damages finally awarded against Licensee directly attributable to any such claim, but only on condition that (a) F5 is notified in writing of such claim within ten days following receipt by Licensee; (b) F5 has sole control of the defense and settlement negotiations, (c) Licensee provides F5 all information and communications received by Licensee concerning such claim, and (d) Licensee provides reasonable assistance to F5 when requested. F5 will have the right, at its option and expense, (i) to obtain for Licensee rights to use the Software, (ii) to replace or modify the Software so it becomes non-infringing , or (iii) to accept return of the Software in exchange or for a credit not to exceed the purchase price paid by Licensee for such Software. The foregoing, subject to the following restrictions, states the exclusive liability of F5 to Licensee concerning infringement.

   a. Restrictions. F5 will have no liability for any claim of infringement based on: (i) use of a superseded or altered release of the Software, (ii) use of the Software in combination with equipment or software

not supplied or specified by F5 in the Software documentation where the Software would not itself be infringing, (iii) use of the Software in an application or environment not described in the Software Documentation or (iv) Software that has been altered or modified in any way by anyone other than F5 or according to F5's instructions.

7. U.S. Government Restricted Rights. The Software was developed at private expense and is provided with "RESTRICTED RIGHTS." Use, duplication or disclosure by the government is subject to restrictions as set forth in FAR 52.227-14 and DFARS 252.227-7013 et. seq. or its successor. The use of this Software by the government constitutes acknowledgment of F5's and its licensors' rights in the Software.

8. DISCLAIMER; LIMITATION OF REMEDY. EXCEPT FOR THE WARRANTIES SPECIFICALLY DESCRIBED HEREIN, F5 AND ITS THIRD PARTY LICENSORS DO NOT MAKE ANY GUARANTEE OR WARRANTY, EXPRESS OR IMPLIED, WITH RESPECT TO THE SOFTWARE, SPECIFICATIONS, SUPPORT, SERVICE OR ANYTHING ELSE. NEITHER F5 NOR ITS THIRD PARTY LICENSORS HAVE AUTHORIZED ANYONE TO MAKE ANY REPRESENTATIONS OR WARRANTIES OTHER THAN AS PROVIDED ABOVE. F5 AND ITS THIRD PARTY LICENSORS DISCLAIM ANY AND ALL WARRANTIES AND GUARANTEES, EXPRESS, IMPLIED OR OTHERWISE, ARISING, WITH RESPECT TO THE SOFTWARE OR SERVICES DELIVERED HEREUNDER, INCLUDING BUT NOT LIMITED TO THE WARRANTY OF MERCHANTABILITY, THE WARRANTY OF FITNESS FOR A PARTICULAR PURPOSE, AND ANY WARRANTY OF NON-INFRINGEMENT OF THE INTELLECTUAL PROPERTY RIGHTS OF ANY THIRD PARTY. THE COLLECTIVE LIABILITY OF F5 AND ITS THIRD PARTY LICENSORS UNDER THIS LICENSE WILL BE LIMITED TO THE AMOUNT PAID FOR THE PRODUCT. F5 AND ITS THIRD PARTY LICENSORS WILL HAVE NO OBLIGATION OR LIABILITY, WHETHER ARISING IN CONTRACT (INCLUDING WARRANTY), TORT (INCLUDING ACTIVE, PASSIVE OR IMPUTED NEGLIGENCE, STRICT LIABILITY OR PRODUCT LIABILITY) OR OTHERWISE FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL OR INDIRECT DAMAGES INCLUDING BUT NOT LIMITED TO LOSS OF USE, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF REVENUE, LOSS OF BUSINESS OR OTHER FINANCIAL LOSS ARISING OUT OF OR IN CONNECTION WITH THE SOFTWARE OR OTHER GOODS OR SERVICES FURNISHED TO LICENSEE BY F5, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

9. Termination. The license granted in Section 2 is effective until terminated, and will automatically terminate if Licensee fails to comply with any of its provisions. Upon termination, Licensee will destroy the Software and documentation and all copies or portions thereof.

10. Miscellaneous. This Agreement will be governed by the laws of the State of Washington, USA without regard to its choice of law rules. The provisions of the U.N. Convention for the International Sale of Goods will not apply. Any provisions found to be unenforceable will not affect the enforceability of the other provisions contained herein, but will instead be replaced with a provision as similar in meaning to the original as possible. This Agreement constitutes the entire agreement between the parties with regard to its subject matter. No modification will be binding unless in writing and signed by the parties.

Table of Contents

Introduction

Getting started I-1
Choosing a configuration tool I-1
Using the Administrator Kit I-2
Stylistic conventions I-3
Finding additional help and technical support resources I-5
What's new in version 3.3 . . . . . I-6
BIG-IP e-Commerce Controller I-6
BIG-IP Cache Controller I-6
Performance enhancements I-7
Learning more about the BIG-IP Controller product family I-7

Chapter 1

Setting up the Hardware

Unpacking and installing the hardware 1-1
Reviewing the hardware requirements 1-1
Familiarizing yourself with the BIG-IP Controller hardware . . . 1-3
Environmental requirements 1-6
Installing and connecting the hardware 1-8

Chapter 2

Working with the First-Time Boot Utility

The First-Time Boot utility 2-1
Gathering the information 2-1
Starting the First-Time Boot utility 2-2
Defining a root password 2-3
Defining a host name 2-3
Configuring a default route 2-4
Configuring a time zone 2-4
Configuring the DNS forwarding proxy settings 2-4

Configuring the interfaces 2-4
Selecting a unit ID 2-5
Choosing a fail-over IP address 2-5
Configuring an interface for the external network 2-7
Configuring an interface for the internal network 2-8
Configuring remote administration 2-10
Configuring settings for the BIG-IP web server 2-11
Confirming your configuration settings 2-12
Committing your configuration settings to the system 2-13

Chapter 3
Additional Setup Options

Introduction to additional setup options 3-1
Defining additional host names 3-1
Preparing workstations for command line access 3-3
Downloading the F-Secure SSH client from the BIG-IP web server 3-4
Downloading the F-Secure SSH client using FTP 3-4
Setting up the F-Secure SSH client on a Windows 95 or Windows NT workstation 3-6
Setting up the F-Secure SSH client on a UNIX workstation . . . 3-7
Addressing general networking issues 3-7
Addressing routing issues 3-8
Configuring DNS on the BIG-IP Controller 3-12
Configuring Email 3-15
Using a serial terminal with the BIG-IP Controller 3-17
To configure a serial terminal in addition to the console 3-19
To configure a serial terminal as the console 3-19
To force a serial terminal to be the console 3-20
Configuring RADIUS authentication 3-21
RADIUS ports on the BIG-IP Controller 3-21
Configuring sshd version 1.3.7 3-22
Configuring sshd version 2.0.12.1 3-23

Index

Introduction

- Getting started
- Using the Administrator Kit
- What's new in version 3.3
- Learning more about the BIG-IP Controller product family

Getting started

Before you start installing the controller, we recommend that you browse the Administrator Guide and find the load balancing solution that most closely addresses your needs. Briefly review the basic configuration tasks and the few pieces of information you should gather in preparation for completing the tasks, such as IP addresses and host names.

Once you find your solution and gather the necessary network information, turn to this Installation Guide for hardware installation instructions, and then return to the Administrator Guide to follow the steps for setting up your chosen solution.

Choosing a configuration tool

The BIG-IP platform offers both web-based and command line configuration tools, so that users can work in the environment that they are most comfortable with.

The First-Time Boot utility

All users will use the First-Time Boot utility, a wizard that walks you through the initial system set up. The First-Time Boot utility automatically starts the first time you turn the controller on, and it prompts you to enter basic system information including a root password and the IP addresses that will be assigned to the network interfaces. The Installation Guide provides detailed information about the specific pieces of information that the First-Time Boot utility prompts you to enter.

The Configuration utility

The Configuration utility is a web-based administrative application that you use to configure and monitor the load balancing setup on the BIG-IP Controller. Once you complete the installation instructions described in this guide, you can use the Configuration utility to the configuration steps outlined in the Administrator Guide. In the Configuration utility, you can also monitor current system performance, and download administrative tools such as the

.

SNMP MIB or the SSH client. The Configuration utility requires Netscape Navigator version 4.7 or later, or Microsoft Internet Explorer version 4.1 or later.

The bigpipe and bigtop command line utilities

The bigpipe™ utility is the command line counter-part to the Configuration utility. Using **bigpipe** commands, you can configure virtual servers, open ports to network traffic, and configure a wide variety of features. To monitor the BIG-IP Controller, you can use certain **bigpipe** commands, or you can use the bigtop™ utility, which provides real-time system monitoring. You can use the command line utilities directly on the BIG-IP Controller, or you can execute commands via a remote shell, such as the SSH client (included with the global release only), or a Telnet client (for countries restricted by cryptography export laws). The BIG-IP Controller Reference Guide provides detailed information about command line syntax.

Using the Administrator Kit

The *BIG-IP® Controller Administrator Kit* provides simple steps for quick, basic configuration, and also provides detailed information about more advanced features and tools, such as the **bigpipe** command line utility. The information is organized into the guides described below.

- ***Installation Guide***
  The Installation Guide walks you through the basic steps needed to get the hardware plugged in and the system connected to the network. Most users turn to this guide only the first time that they set up a BIG-IP Controller. The Installation Guide also covers general network administration issues, such as setting up common network administration tools including Sendmail.
- ***Administrator Guide***
  The Administrator Guide provides examples of common load balancing solutions supported by the particular type of BIG-IP Controller you purchased. For example, in the BIG-IP HA

Controller Administrator Guide, you can find everything from a basic web server load balancing solution to a firewall load balancing solution.

- ***Reference Guide***
  The Reference Guide provides basic descriptions of individual BIG-IP objects, such as pools, nodes, and virtual servers. It also provides syntax information for **bigpipe** commands, configuration utilities, configuration files, and system utilities.
- ***F-Secure SSH User Guide***
  This guide is distributed only with BIG-IP Controllers that support the F-Secure SSH client (a tool used for remote command line access). It provides information about setting up and using the SSH client.

Stylistic conventions

To help you easily identify and understand certain types of information, all F5 Networks administrative documentation uses the stylistic conventions described below.

WARNING

All examples in F5 Networks documentation use only non-routable IP addresses. When you set up the solutions we describe, you must use IP addresses suitable to your own network in place of our sample addresses.

Identifying new terms

When we first define a new term, the term is shown in bold italic text. For example, a ***virtual server*** is a the combination of an IP address and port that maps to a set of back-end servers.

.

Identifying references to objects, names, and commands

We apply bold text to a variety of items to help you easily pick them out of a block of text. These items include web addresses, IP addresses, utility names, and portions of commands, such as variables and keywords. For example, the **bigpipe vip** command requires that you include at least one <**node**> variable.

Identifying references to other documents

We use italic text to denote a reference to another document. In references where we provide the name of a book as well as a specific chapter or section in the book, we show the book name in bold, italic text, and the chapter/section name in italic text to help quickly differentiate the two. For example, you can find information about bigpipe commands in the *bigpipe Command Referenc*e section of the ***BIG-IP Controller Reference Guide***.

Identifying command syntax

We show actual, complete commands in bold Courier text. Note that we do not include the corresponding screen prompt, unless the command is shown in a figure that depicts an entire command line screen. For example, the following command sets the BIG-IP Controller load balancing mode to Round Robin:

```
bigpipe lb rr
```

Table 1 explains additional special conventions used in command line syntax.

Item in text	Description
****	Continue to the next line without typing a line break.
< >	You enter text for the enclosed item. For example, if the command has **<your name>**, type in your name.
|	Separates parts of a command.
[ ]	Syntax inside the square brackets is optional.
...	Indicates that you can type a series of items.

Table 1 *Command line syntax conventions*

Finding additional help and technical support resources

In addition to this administrator guide, you can find technical documentation about the BIG-IP Controller in the following locations:

- **Release notes**
  The release note for the current version of the BIG-IP Controller is available from the web server on the BIG-IP Controller. The release note contains the latest information for the current version, including a list of new features and enhancements, a list of fixes, and, in some cases, a list of known issues.
- **Online help for BIG-IP Controller features**
  You can find help online in three different locations:
  - The web server on the BIG-IP Controller has PDF versions of the guides included in the Administrator Kit. BIG-IP Controller upgrades replace these guides with updated versions as appropriate.
  - The web-based Configuration utility has online help for each screen. Simply click the **Help** button in the toolbar.
  - Individual **bigpipe** commands have online help, including command syntax and examples, in standard UNIX man page format. Simply type the command followed by the question mark option (**-?**), and the BIG-IP Controller displays the syntax and usage associated with the command.
- **Third-party documentation for software add-ons**
  The web server on the BIG-IP Controller contains online documentation for all third-party software included with the BIG-IP Controller, such as GateD.
- **Technical support via the World Wide Web**
  The F5 Networks Technical Support web site, **http://tech.F5.com**, provides the latest technical notes, answers to frequently asked questions, updates for administrator guides (in PDF format), and the Ask F5 natural language question and answer engine. To access this site, you need to obtain a customer ID and a password from the F5 Help Desk.

.

What's new in version 3.3

The BIG-IP Controller offers the following major new features in version 3.3, in addition to many smaller enhancements.

BIG-IP e-Commerce Controller

The BIG-IP e-Commerce Controller is a new member of the BIG-IP product family. You can use the BIG-IP e-Commerce Controller to process SSL connections to your network. This controller contains a specific set of software and hardware features that accelerate SSL connections.

BIG-IP Cache Controller

This version of the BIG-IP Controller is available as the BIG-IP Cache Controller. The BIG-IP Cache Controller version contains a specific set of features from the BIG-IP Controller that maximizes the efficiency of caches in your network. In addition to the load balancing features available with this controller, this version of the controller has new rule syntax that provides the ability to redirect HTTP requests to caches in your network. These features include:

- **Cacheable content determination**
  This feature enables you to determine the type of content you cache on the basis of any combination of elements in the header of an HTTP request.
- **Content affinity**
  This feature assures that the same cache serves the same content subset even when caches become temporarily unavailable or when caches are added to or deleted from the cache pool.
- **Hot content load balancing**
  When configured, this feature identifies highly requested content and redirects these requests to a hot pool for load balancing.
- **Intelligent cache population**
  When configured, this feature allows caches to retrieve content from other caches in addition to the origin web server.

Performance enhancements

This version of the BIG-IP Controller includes internal performance enhancements. These enhancements improve the overall performance of the BIG-IP Controller.

Learning more about the BIG-IP Controller product family

The BIG-IP Controller platform offers many different software systems. These systems can be stand-alone, or can run in redundant pairs, with the exception of the BIG-IP e-Commerce Controller, which is only available as a stand-alone system. You can easily upgrade from any special-purpose BIG-IP Controller to the BIG-IP HA Controller, which supports all BIG-IP Controller features.

- **The BIG-IP LB Controller**
  The BIG-IP LB Controller provides basic load balancing features.
- **The BIG-IP FireGuardController**
  The BIG-IP FireGuard Controller provides load balancing features that maximize the efficiency and performance of a group of firewalls.
- **The BIG-IP Cache Controller**
  The BIG-IP Cache Controller uses content-aware traffic direction to maximize the efficiency and performance of an group of cache servers.
- **The BIG-IP e-Commerce Controller**
  The BIG-IP e-Commerce Controller uses SSL acceleration technology to increase the speed and reliability of the secure connections that drive e-commerce sites.

❖ **The BIG-IP HA Controller**
The BIG-IP HA Controller provides all features from the basic BIG-IP LB Controller to the advanced BIG-IP FireGuard, BIG-IP Cache Controller, and BIG-IP e-Commerce Controller products.

◆ Note

BIG-IP Controllers distributed outside of the United States to a select few countries, regardless of system type, do not support encrypted communications. They do not include the F-Secure SSH client, nor do they support SSL connections to the BIG-IP web server. Instead, you can use the standard Telnet, FTP, and HTTP protocols to connect to the unit and perform administrative functions.

F5 Networks BIG-IP Local Traffic Controller

1

Setting Up the Hardware

- Unpacking and installing the hardward

Unpacking and installing the hardware

There are four basic tasks you must complete to get the BIG-IP Controller installed and set up.

- Review the hardware requirements
- Familiarize yourself with the BIG-IP Controller hardware
- Review the environmental requirements
- Connect the peripheral hardware and connect the BIG-IP Controller to the network

Reviewing the hardware requirements

The BIG-IP Controller comes with the hardware that you need for installation and maintenance. However, you must provide standard peripheral hardware, such as a keyboard or serial terminal.

Hardware provided with the BIG-IP Controller

When you unpack the BIG-IP Controller, you should make sure that the following components are included:

- One power cable
- One PC/AT-to-PS/2 keyboard adapter
- Four rack-mounting screws
- Two keys for the front panel lock
- One extra fan filter
- One BIG-IP documentation kit
  The BIG-IP documentation kit contains the following items:
  - *The BIG-IP Controller Installation Guide*
  - *The BIG-IP Controller Administrator Guide*
  - *The BIG-IP Controller Reference Guide*
  - Additionally, if you purchase the global release of the BIG-IP Controller that supports encryption, you receive the ***F-Secure SSH Client*** manual, published by Data Fellows.

If you purchased a hardware-based redundant system, you also received one fail-over cable to connect the two controller units together (network-based redundant systems do not require a fail-over cable).

Peripheral hardware that you provide

For each BIG-IP Controller in the system, you need to provide the following peripheral hardware:

- You need standard input/output hardware for direct administrative access to the BIG-IP Controller. Either of the following options is acceptable:
  - A VGA monitor and PC/AT-compatible keyboard
  - Optionally, a serial terminal and a null modem cable (see ***Using a serial terminal with the BIG-IP Controller***, on page 3-17, for serial terminal configuration information)
- You also need network hubs, switches, or concentrators to connect to the BIG-IP Controller network interfaces. The devices you select must be compatible with the network interface cards installed in the BIG-IP Controller. The devices can support 10/100 Ethernet, Gigabit Ethernet, or FDDI/CDDI (including multiple FDDI and full duplex).
  - For Ethernet, you need either a 10Mb/sec or 100 Mb/sec hub or switch
  - For FDDI/CDDI, a concentrator or a switch is optional

If you plan on doing remote administration from your own PC workstation as most users do, we recommend that you have your workstation already in place. Keep in mind that the First-Time Boot utility prompts you to enter your workstation's IP address when you set up remote administrative access.

Familiarizing yourself with the BIG-IP Controller hardware

The BIG-IP Controller is offered in 4U and 2U hardware configurations. Before you begin to install the BIG-IP Controller, you may want to quickly review the following figures that illustrate the controls and ports on both the front and the back of a 4U BIG-IP Controller and a 2U BIG-IP Controller.

Using the BIG-IP Controller 4U hardware configuration

This section describes the front and back layout of a 4U BIG-IP Controller. If you have a special hardware configuration, such as those that include more than two interface cards, the ports on the back of your unit will differ slightly from those shown below.

◆ Note

The ports on the back of every BIG-IP Controller are individually labeled.

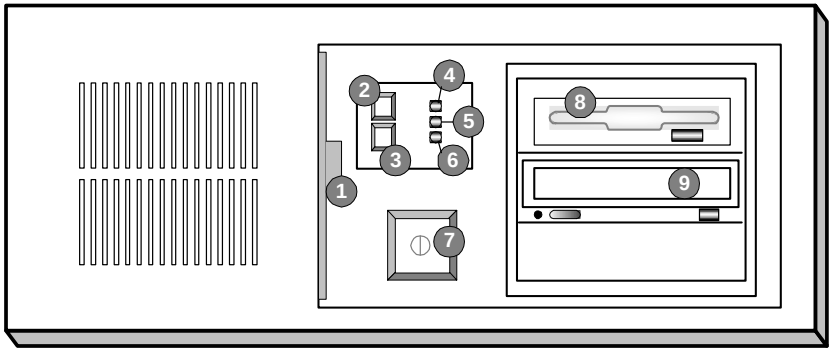

1. Fan filter	6. Power LED
2. Keyboard lock	7. On/off button
3. Reset button	8. 3.5 floppy disk drive
4. Keyboard lock LED	9. CD-ROM drive
5. Hard disk drive LED	

Figure 1.1 *Front view of a 4U BIG-IP Controller*

Figure 1.1 illustrates the front of a 4U BIG-IP Controller with the access panel open. On the front of the unit, you can turn the unit off and on, or you can reset the unit. You can also view the indicator lights for hard disk access and for the keyboard lock.

Figure 1.2, the following figure, illustrates the back of a 4U BIG-IP Controller. Note that all ports are labeled, even those which are not intended to be used with the BIG-IP Controller. Ports marked with an asterisk (*) in the list following are not used by the BIG-IP Controller, and do not need to be connected to any peripheral hardware.

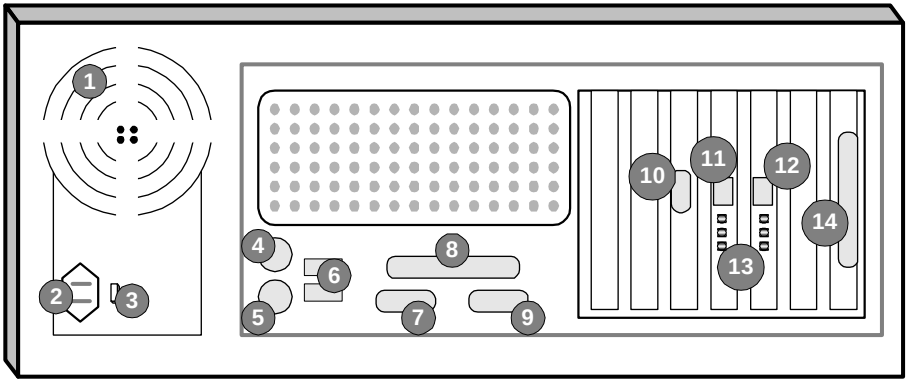

1. Fan	8. Printer port*
2. Power in	9. Fail-over port
3. Voltage selector	10. Video (VGA) port
4. Mouse port*	11. Internal interface (RJ-45)
5. Keyboard port	12. External interface (RJ-45)
6. Universal serial bus ports*	13. Interface indicator LEDs
7. Serial terminal port	14. Watchdog card*

**Not to be connected to any peripheral hardware.*

Figure 1.2 *Back view of a 4U BIG-IP Controller*

Using the BIG-IP Controller 2U hardware configuration

This section describes the front and back layout of a 2U BIG-IP Controller. If you have a special hardware configuration, such as those that include more than two interface cards, the ports on the back of your unit will differ slightly from those shown below.

◆ Note

The ports on the back of every BIG-IP Controller are individually labeled, so it should be clear what each port is, no matter which hardware configuration you have purchased.

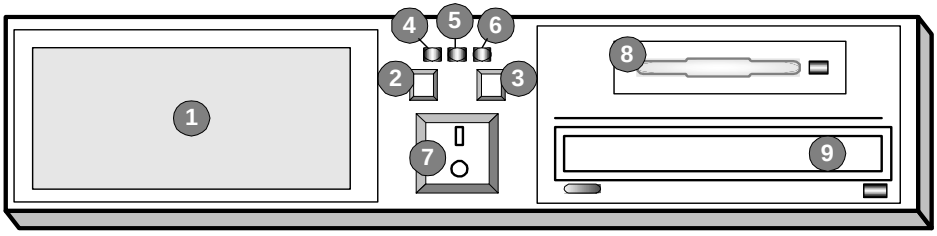

1. Fan filter	6. Power LED
2. Keyboard lock	7. On/off button
3. Reset button	8. Flash or PC card
4. Keyboard lock LED	9. CD-ROM drive
5. Hard disk drive LED	

Figure 1.3 *Front view of a 2U BIG-IP Controller*

Figure 1.3 illustrates the front of a 2U BIG-IP Controller with the access panel open. On the front of the unit, you can turn the unit off and on, or you can reset the unit. You can also view the indicator lights for hard disk access and for the keyboard lock.

Figure 1.4, the following figure, illustrates the back of a 2U BIG-IP Controller. Note that all ports are labeled, even those which are not intended to be used with the BIG-IP Controller. Ports marked with

an asterisk (*) in the list following are not used by the BIG-IP Controller, and do not need to be connected to any peripheral hardware.

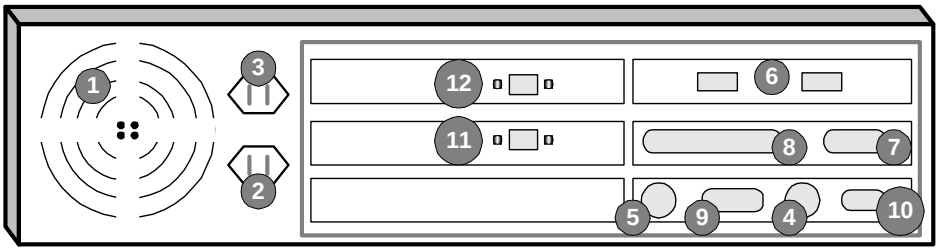

1. Fan
2. Power cable plug
3. Not used
4. Mouse port*
5. Keyboard port
6. Universal serial bus ports*
7. Serial terminal port
8. Printer port*
9. Fail-over port
10. Video (VGA) port
11. Internal interface (RJ-45)
12. External interface (RJ-45)

**Not to be connected to any peripheral hardware.*

Figure 1.4 *Back view of a 2U BIG-IP Controller*

Environmental requirements

General guidelines

A BIG-IP Controller is an industrial network appliance, designed to be mounted in a standard 19-inch rack. To ensure safe installation and operation of the unit:

- Install the rack according to the manufacturer's instructions, and check the rack for stability before placing equipment in it.
- Build and position the rack so that once you install the BIG-IP Controller, the power supply and the vents on both the front and back of the unit remain unobstructed. The BIG-IP Controller must have adequate ventilation around the unit at all times.
- Do not allow the air temperature in the room to exceed 40° C.

- Do not plug the unit into a branch circuit shared by more electronic equipment than the circuit is designed to manage safely at one time.
- Verify that the voltage selector is set appropriately before connecting the power cable to the unit.

The unit must be connected to Earth ground, and it should have a reliable ground path maintained at all times.

The BIG-IP Controller contains a lithium battery. There is danger of an explosion if you replace the lithium battery incorrectly. We recommend that you replace the battery only with the same type of battery originally installed in the unit, or with an equivalent type recommended by the battery manufacturer. Be sure to discard all used batteries according to the manufacturer's instructions.

This equipment is not intended for operator serviceability. To prevent injury and to preserve the manufacturer's warranty, allow only qualified service personnel to service the equipment.

Guidelines for DC powered equipment

A DC powered installation must meet the following requirements:

- Install the unit using a 20 Amp external branch circuit protection device.

- For permanently connected equipment, incorporate a readily accessible disconnect in the fixed wiring.
- Use only copper conductors.

Install DC powered equipment only in restricted access areas, such as dedicated equipment rooms, equipment closets, or similar locations.

Installing and connecting the hardware

There are six basic steps to installing the hardware. You simply need to install the controller in the rack, connect the peripheral hardware and the external and internal interfaces, and then connect the fail-over and power cables. If you have a unit with three or more network interface cards (NICs), be sure to review step 3.

WARNING

Do not turn on a BIG-IP Controller until all peripheral hardware is connected to the unit.

To install the hardware

1. Insert the BIG-IP Controller in the rack and secure it using the four rack-mounting screws that are provided.
2. Connect the hardware that you have chosen to use for input/output:
   - If you are using a VGA monitor and keyboard, connect the monitor connector cable to the video port (number 10 in Figure 1.2 for 4U, or in Figure 2.4 for 2U) and the keyboard connector cable to the keyboard port (number 5 in Figure 1.2 for 4U, or in Figure 2.4 for 2U). Note that a PC/AT-to-PS/2 keyboard adapter is included with each BIG-IP Controller (see the component list on page 1-1).

- Optionally, if you are using a serial terminal as the console, connect the serial cable to the terminal serial port (number 7 in Figure 1.2 for 4U, or in Figure 2.4 for 2U). Also, you should not connect a keyboard to the BIG-IP Controller. If there is no keyboard connected to the BIG-IP Controller when it is started or rebooted, the BIG-IP Controller defaults to using the serial port as the console.

3. Connect the external interface (number 12 in Figure 1.2 for 4U, or in Figure 2.4 for 2U) to the network from which the BIG-IP Controller receives connection requests.

   - If you have purchased a unit with three or more network interface cards (NICs), be sure to note or write down how you connect the cables to the internal and external interfaces. When you run the First-Time Boot utility, it automatically detects the number of interfaces that are installed and prompts you to configure more external interfaces, if you want. It is important to select the correct external interface based on the way you have connected the cables to the back of the unit.

4. Connect the internal interface (number 11 in Figure 1.2 for 4U, or in Figure 2.4 for 2U) to the network that houses the array of servers, routers, or firewalls that the BIG-IP Controller load balances.

5. If you have a hardware-based redundant system, connect the fail-over cable to the terminal serial port on each unit (number 7 in Figure 1.2 for 4U, or number 7 in Figure 2.4 for 2U).

6. Connect the power cable to the BIG-IP Controller (number 2 in Figure 1.2 for 4U, or Figure 2.4 for 2U), and then connect it to the power source.

WARNING

Before connecting the power cable to a power supply, customers outside the United States should make sure that the voltage selector is set appropriately. This check is necessary only if the controller has an external voltage selector.

2

Working with the First-Time Boot Utility

- The First-Time Boot utility
- Gathering the information
- Starting the First-Time Boot utility
- Defining a root password
- Defining a host name
- Configuring a default route
- Configuring a time zone
- Configuring DNS forwarding proxy settings
- Configuring the interfaces
- Configuring remote administration
- Configuring settings for the BIG-IP web server
- Confirming your configuration settings
- Committing your configuration settings to the system

The First-Time Boot utility

The First-Time Boot utility is a wizard that walks you through a brief series of required configuration tasks, such as defining a root password, and configuring IP addresses for the interfaces. Once you complete the First-Time Boot utility, you can connect to the BIG-IP Controller from a remote workstation and begin configuring your load balancing setup.

The First-Time Boot utility is organized into three phases: configure, confirm, and commit. Each phase guides you through a series of screens, presenting the information in the following order:

- Root password
- Host name
- Default route (typically a router's IP address)
- Time zone
- DNS forwarding proxy
- Interface settings for each network interface
- Configuration for BIG-IP Controller redundant systems (fail-over IP address)
- IP address for remote administration
- Settings for the web server on the BIG-IP Controller

First, you configure all of the required information. Next, you have the opportunity to confirm each individual setting or correct it if necessary. Last, your confirmed settings are committed and saved to the system. Note that the screens you see are tailored to the specific hardware and software configuration that you have. If you have a stand-alone system, for example, the First-Time Boot utility skips the redundant system screens.

Gathering the information

Before you run the First-Time Boot utility on a specific BIG-IP Controller, you should have the following information ready to enter:

- Passwords for the terminal login and for the BIG-IP web server
- Host names for the terminal login and for the BIG-IP web server
- A default route (typically a router's IP address)
- Settings for the network interfaces, including IP addresses, media type, and optionally a custom netmask and broadcast addresses
- Configuration information for redundant systems, including an IP alias for the shared address, and the IP address of the corresponding unit
- The IP address or IP address range for remote administrative connections

Starting the First-Time Boot utility

The First-Time Boot utility starts automatically when you turn on the BIG-IP Controller (the power switch is located on the front of the BIG-IP Controller). The first screen the BIG-IP Controller displays is the License Agreement screen. You must scroll through the screen, read it, and accept the agreement before you can move to the next screen. If you agree to the license statement, the next screen you see is the Welcome screen. From this screen, simply press any key on the keyboard to start the First-Time Boot utility, and then follow the instructions on the subsequent screens to complete the process.

◆ Note

You can re-run the First-Time Boot utility after you run it for initial configuration. To re-run the First-Time Boot utility, type ***config*** *on the command line.*

Defining a root password

A root password allows you command line administrative access to the BIG-IP Controller system. The password must contain a minimum of 6 characters, but no more than 32 characters. Passwords are case-sensitive, and we recommend that your password contain a combination of upper- and lower-case characters, as well as numbers and punctuation characters. Once you enter a password, the First-Time Boot utility prompts you to confirm your root password by typing it again. If the two passwords match, your password is immediately saved. If the two passwords do not match, the First-Time Boot utility provides an error message and prompts you to re-enter your password.

WARNING

*The root password is the only setting that is saved immediately, rather than confirmed and committed at the end of the First-Time Boot utility process. You cannot change the root password until the First-Time Boot utility completes and you reboot the BIG-IP Controller (see the **BIG-IP Controller Administration Guide**, Monitoring and Administration). Note that you can change other system settings when the First-Time Boot utility prompts you to confirm your configuration settings.*

Defining a host name

The host name identifies the BIG-IP Controller itself. Host names must start with a letter, and must be at least two characters. They may contain numbers, letters, and the symbol for dash (-). There are no additional restrictions on host names, other than those imposed by your own network requirements.

Configuring a default route

If a BIG-IP Controller does not have a predefined route for network traffic, the controller automatically sends traffic to the IP address that you define as the default route. Typically, a default route is set to a router's IP address.

Configuring a time zone

Next, you need to specify your time zone. This ensures that the clock for the BIG-IP Controller is set correctly, and that dates and times recorded in log files correspond to the time zone of the system administrator. Scroll through the list to find the time zone at your location. Note that one option may appear with multiple names. Select the time zone you want to use, and press the Enter key to continue.

Configuring the DNS forwarding proxy settings

You only need to complete this step if you want machines inside your BIG-IP managed network to use DNS servers outside of that network (e.g., for reverse DNS lookup from a web server).

Specify the DNS name server and domain name for DNS proxy forwarding by the BIG-IP Controller.

Configuring the interfaces

On the Configure BIG-IP Interfaces screen, select **Yes** if you have a redundant system.

Selecting a unit ID

If you are configuring a redundant system, the First-Time Boot utility prompts you to provide a unit ID and the IP address for fail-over for the BIG-IP Controller. The default unit ID number is **1**. If this is the first controller in the redundant system, use the default. When you configure the second controller in the system, type **2**. These unit IDs are used for active-active redundant controller configuration.

Choosing a fail-over IP address

If you are configuring a redundant system, after you type in a unit number, the First-Time Boot utility prompts you to provide an IP address for fail-over. Type in the IP address configured on the internal interface of the other BIG-IP Controller.

Configuring internal and external interfaces

We recommend that you configure at least one external interface, and at least one internal interface on each controller. The external interface is the one on which the BIG-IP Controller receives connection requests. The internal interface is the one that is connected to the network of servers, firewalls, or other equipment that the BIG-IP Controller load balances. The utility prompts you for each interface, and asks you to provide the IP address, netmask, broadcast address, and the interface media type. With this release of the BIG-IP Controller, the concept of interfaces as internal and external is changing. You can now choose each attribute you want to assign to an interface. In effect, this means that you can

configure one interface with the properties of both an internal and external interface. Table 2.1 shows the attributes that determine the way an interface handles connections.

Interface type	Attributes
Internal	Process source addresses Administrative ports open
External	Process destination addresses Administrative ports locked

Table 2.1 *Attributes of internal and external interfaces*

◆ Note

After you complete the First-Time Boot utility, you can change the individual attributes of an interface. For information about changing interface attributes, see the ***Reference Guide****.*

If you have a redundant system, the First-Time Boot utility prompts you to provide the IP address that serves as an alias for both BIG-IP Controllers. The IP alias is shared between the units, and is used by active controllers. Each unit also uses unique internal and external IP addresses. The First-Time Boot utility guides you through configuring the interfaces, based on your hardware configuration.

We recommend that you set the internal alias as the default route for the node servers. Note that for each IP address or alias that you assign to an interface, you have the option of assigning a custom netmask and broadcast address as well.

Configuring an interface for the external network

The Select Interfaces screen shows a list of the installed interfaces. Select the one you want to use for the external network, and press the Enter key.

Note

The IP address of the external network interface is not the IP address of your site or sites. The IP addresses of the sites themselves are specified by the virtual IP addresses associated with each virtual server you configure.

WARNING

The configuration utility lists only the network interface devices that it detects during boot up. If the utility lists only one interface device, the network adapter may have come loose during shipping. Check the LED indicators on the network adapters to ensure that they are working and are connected.

Once you select the interface, the utility prompts you for the following information, in many cases offering you a default:

- **IP address**
- **Netmask**
  Note that the BIG-IP Controller uses a default netmask appropriate to the subnetwork indicated by the IP address.
- **Broadcast address**
  The default broadcast address is a combination of the IP address and the netmask.
- **Shared IP alias (redundant systems only)**
  The external IP alias associated with each unit's external interface
- **Shared IP alias netmask (redundant systems only)**
- **Shared IP alias broadcast address (redundant systems only)**

❖ **Media type for Interface**
The media type options depend on the network interface card included in your hardware configuration. The BIG-IP Controller supports the following types:

- auto
- 10baseT
- 10baseT,FDX
- 100baseTX
- 100baseTX,FDX
- FDDI
- Gigabit Ethernet

If you are configuring a BIG-IP Controller that has more than two network interface cards installed, the First-Time Boot utility prompts you to configure more external interfaces. When you complete the configuration of an interface, you return to the Interface Configuration screen and repeat the steps described above.

◆ Tip

We recommend that you configure at least one internal interface.

Configuring an interface for the internal network

When you configure the interface that connects the BIG-IP Controller to the internal network (the servers and other network devices that sit behind the BIG-IP Controller), the First-Time Boot utility prompts you for the following information:

❖ **IP address**

❖ **Netmask**
Note that the BIG-IP Controller uses a default netmask appropriate to the subnetwork indicated by the IP address.

❖ **Broadcast address**
The default broadcast address is a combination of the IP address and the netmask.

- **Shared IP alias (redundant systems only)**
  An IP alias associated with each unit's internal interface
- **Shared IP alias netmask (redundant systems only)**
- **Shared IP alias broadcast address (redundant systems only)**
- **Media type for Interface**
  The media type options depend on the network interface card included in your hardware configuration. The BIG-IP Controller supports the following types:
  - auto
  - 10baseT
  - 10baseT,FDX
  - 100baseTX
  - 100baseTX,FDX
  - FDDI
  - Gigabit Ethernet

◆ Note

We recommend that you set the default route of each network device behind the BIG-IP redundant system to the internal IP alias of the BIG-IP Controllers. This guarantees that the network devices always communicate with an active BIG-IP Controller in the redundant system.

If you configure more than one internal interface on a redundant system, the First-Time Boot utility prompts you to choose one as the primary internal interface. The interface you choose as the primary internal interface is used for exchanging network based fail-over and state fail-over information with the other controller in a redundant system.

Configuring remote administration

On most BIG-IP Controllers, the first screen you see is the Configure SSH screen, which prompts you to type an IP address for SSH command line access. If SSH is not available, you are prompted to configure access through Telnet and FTP instead.

When you configure shell access, the First-Time Boot utility prompts you to create a support account for that method. You can use this support account to provide an F5 Networks engineer access to the BIG-IP Controller.

When the First-Time Boot utility prompts you to enter an IP address for administration, you can type a single IP address or a range of IP addresses, from which the BIG-IP Controller will accept administrative connections (either remote shell connections, or connections to the BIG-IP web server). To specify a range of IP addresses, you can use the asterisk (*****) as a wildcard character in the IP addresses.

The following example allows remote administration from all hosts on the 192.168.2 network:

```
192.168.2.*
```

◆ Tip

In order to use the configuration synchronization feature for redundant units, you must configure the BIG-IP Controller for command line access.

◆ Note

For administration purposes, you can connect to the BIG-IP Controller IP alias, which always connects you to an active controller. To connect to a specific controller, simply connect directly to the IP address of that BIG-IP Controller.

Configuring settings for the BIG-IP web server

The BIG-IP web server requires you to define a fully qualified domain name (FQDN) for the server on each interface. The BIG-IP web server configuration also requires that you define a user ID and password. If SSL is available, the configuration also generates authentication certificates.

The First-Time Boot utility guides you through a series of screens to set up web server access.

- The first screen prompts you to select the interface you want to configure for web access. After you select an interface to configure, the utility prompts you to type a fully qualified domain name (FQDN) for the interface. You can configure web access on one or more interfaces.
- After you configure the interface, the utility prompts you for a user name and password. After you type a user name and password, the utility prompts you for a vendor support account. The vendor support account is not required.
- The certification screen prompts you for country, state, city, company, and division.

Once you have completed this screen, the First-Time Boot utility moves into the confirmation phase.

WARNING

If you ever change the IP addresses or host names on the BIG-IP Controller interfaces, you must reconfigure the BIG-IP web server to reflect your new settings. You can run the re-configuration utility from the command line using the following command:

```
reconfig-httpd
```

You can also add users to the existing password file, change a password for an existing user, or recreate the password file, without actually going through the BIG-IP web server configuration process. For more information, see the ***BIG-IP Controller Reference Guide***, *BIG-IP Controller Configuration Utilities*.

WARNING

If you have modified the BIG-IP web server configuration outside of the configuration utility, be aware that some changes may be lost when you run the ***reconfig-httpd*** *utility. This utility overwrites the* ***httpd.conf*** *file, and several other files, but it does warn you before doing so.*

Confirming your configuration settings

At this point, you have entered all the configuration information, and now you simply have to confirm each setting. Each confirmation screen displays a setting, and prompts you to accept or re-enter it. If you choose to edit it, the utility displays the same configuration screen in which you defined the setting the first time. When you finish editing the item, you return directly to the Confirmation screen for that item, and continue the confirmation process. Note that once you accept a setting in the Confirmation screen, you do not have another opportunity to review it.

You confirm or edit the settings in the same order that you configured them:

- Confirm Host name
- Confirm Default route
- Confirm time zone
- Confirm all interface settings
- Confirm fail-over IP address, if necessary
- Confirm administrative IP address
- Confirm web server options

Once you have confirmed the last setting, the First-Time Boot utility moves directly into the commit phase, where you are not able to make any changes.

Committing your configuration settings to the system

Once you confirm all of the configuration settings, the configuration utility saves the configuration settings. During this commit process, the First-Time Boot utility creates the following files and configuration database records:

- An **/etc/hosts.allow** file
  This file stores the IP address, or IP address range, from which the BIG-IP Controller accepts administrative connections.
- Interface entries in BIG/db
- An **/etc/bigip.conf** file
- An **/etc/netstart** file
- An **/etc/hosts** file
- An **/etc/ethers** file
- A **/var/f5/httpd/conf/httpd.conf** file
- An **/etc/sshd_config** file

If you want to change any information in these files at a later time, you can edit the files directly, you can change the information in the web-based Configuration utility, or for certain settings, you can change them using command line utilities. If necessary, you can also re-run the First-Time Boot utility.

3

Additional Setup Options

- Introduction to additional setup options
- Defining additional host names
- Preparing workstations for command line access
- Addressing general networking issues
- Using a serial terminal with the BIG-IP Controller
- Configuring RADIUS authentication

Introduction to additional setup options

This chapter contains details about additional setup options you may want to configure for the controller. The options described in this chapter include:

- Defining additional host names
- Preparing workstations for command line access
- Addressing general networking issues
- Using a serial terminal with the BIG-IP Controller
- Configuring RADIUS authentication

Defining additional host names

Once you complete the First-Time Boot utility, you may want to insert additional host names and IP addresses for network devices into the **/etc/hosts** file to allow for more user-friendly system administration. In particular, you may want to create host names for the IP addresses that you will assign to virtual servers. You may also want to define host names for standard devices such as your routers, network interface cards, and the servers or other equipment that you are load balancing.

The **/etc/hosts** file, as created by the First-Time Boot utility, is similar to the example shown in Figure 3.1.

```
# localhost entry
127.1   localhost

# default gateway entry
11.11.11.10     router


# Local name
11.11.11.2   bigip controller name

#
# Physical Interfaces Tue Oct 19 18:14:44  1999
#

# ext interface
11.11.11.2      exp0

# int interface
11.12.11.2      exp1

#
# VIPS and NODES ( add below - do not delete this line )
#
```

Figure 3.1 *The **/etc/hosts** file created by the First-Time Boot utility*

This sample **hosts** file lists the IP addresses for the default router, the internal network interface, and the external network interface, and it contains place holders for both the virtual servers and the content servers that the BIG-IP Controller will manage.

WARNING

*If you have modified the **/etc/hosts** file with something other than the First-Time Boot utility, such as **vi** or **pico**, be aware that your changes may be lost when you run the First-Time Boot utility (**config file**). This utility overwrites the **/etc/hosts** file, and several other files, but it does warn you before doing so.*

Preparing workstations for command line access

You may want to configure a workstation for command line access to the BIG-IP Controller. You can use a workstation configured for command line access to configure the BIG-IP Controller remotely.

The type of system you have determines the options you have for remote command line administration:

- Most BIG-IP Controllers support secure shell command line access using the F-Secure SSH client.
- If your BIG-IP Controllers does not support SSH client, you will use command line access using a standard Telnet shell.

If SSH came with your BIG-IP Controller, you probably want to install the F-Secure SSH client on your workstation. The BIG-IP Controller includes a version of the F-Secure SSH client for each of the following platforms: Windows, UNIX, and Macintosh. You can download the F-Secure client using your web browser, or you can download the client using an FTP server on the administrative workstation.

Note that the F-Secure license agreement allows you to download two copies of the F-Secure SSH client. If you require additional licenses, you need to contact Data Fellows. For information about

contacting Data Fellows, as well as information about working with the SSH client, refer to the F-Secure manual included with your BIG-IP Controller.

Note

You can also use the F-Secure SSH suite for file transfer to and from the BIG-IP Controller, as well as for remote backups. An F-Secure SSH client is pre-installed on the BIG-IP Controller to assist with file transfer activities. Please refer to the F-Secure User's Manual for more information.

Downloading the F-Secure SSH client from the BIG-IP web server

The F-Secure SSH client is available in the Downloads section of the BIG-IP web server. For US products, you connect to the BIG-IP web server via SSL on port 443 (use **https://** rather than **http://** in the URL). Once you connect to the BIG-IP web server, go to the Additional Software Downloads section and click the SSH Clients link. From the SSH Clients page, you can select the SSH Client.

Downloading the F-Secure SSH client using FTP

The BIG-IP Controller has an FTP client installed, which allows you to transfer the F-Secure SSH Client using FTP (note that your destination workstation must also have an FTP server installed). After you transfer the installation file, you simply decompress the file and run the F-Secure installation program.

Note

*You can allow FTP and Telnet access to the BIG-IP Controller by running the **config_ftpd** script from the command line. This script allows specific clients FTP or Telnet access to the BIG-IP Controller. However, we do not recommend this method. For more information about this script, refer to the **BIG-IP Controller Reference Guide**.*

You can initiate the FTP transfer from the BIG-IP Controller using the attached monitor and keyboard.

To transfer the SSH client using FTP

1. Locate the SSH client that is appropriate for the operating system that runs on the administrative workstation:
   - Change directories to the ***/usr/contrib/fsecure*** directory where the F-secure SSH clients are stored.
   - List the directory, noting the file name that corresponds to the operating system of your administration workstation.
2. To start FTP, type the following command:

   ```
   ftp
   ```
3. To open a connection to the remote workstation, type the following command (where **IP address** is the IP address of the remote workstation itself):

   ```
   open <IP address>
   ```

   Once you connect to the administrative workstation, the FTP server on the administrative workstation prompts you for a password.
4. Enter the appropriate user name and password to complete the connection.
5. To switch to passive FTP mode, type the following command:

   ```
   passive
   ```
6. To switch the transfer mode to binary, type the following command:

   ```
   bin
   ```
7. Go to the directory on the administrative workstation where you want to install the F-Secure SSH client.

8. To start the transfer process, type the following command (where **filename** is the name of the F-Secure file that is specific to the operating system running on the administrative workstation):

   ```
   put <filename>
   ```

9. Once the transfer is done, type the following command:

   ```
   quit
   ```

Setting up the F-Secure SSH client on a Windows 95 or Windows NT workstation

The F-Secure SSH client installation file for Windows platforms is compressed in ZIP format. You can use standard ZIP tools, such as PKZip or WinZip to extract the file.

To unzip and install the SSH client

1. Log on to the Windows workstation.
2. Go to the directory to which you transferred the F-Secure installation file. Run **PKZip** or **WinZip** to extract the files.
3. The set of files extracted includes a Setup program. Run the Setup program to install the client.
4. Start the F-Secure SSH client.
5. In the SSH Client window, from the Edit menu choose **Properties**.
   The Properties dialog box opens.
6. In the Connection tab, in the Remote Host section, type the following items:
   - In the **Host Name** box, type the BIG-IP Controller IP address or host name.
   - In the **User Name** box, type the root user name.
7. In the Options section, check **Compression** and set the Cipher option to **Blowfish**.
8. Click the **OK** button.

Setting up the F-Secure SSH client on a UNIX workstation

The F-Secure installation file for UNIX platforms is compressed in TAR/Gzip format.

To untar and install the SSH client

1. Log on to the workstation and go to the directory into which you transferred the F-Secure SSH client tar file.
2. Untar the file and follow the instructions in the **install** file to build the F-Secure SSH client for your workstation.
3. Start the SSH client.
4. Open a connection to the BIG-IP Controller:

   ```
   ssh -l root [BIG-IP IP address]
   ```

5. Type the root password and press the Enter key.

Addressing general networking issues

You must address several network issues when you place a BIG-IP Controller in your network. These networking issues include routing, DNS configuration, and special e-mail considerations. You need to address these issues based on the type of hardware and software in your network. This section describes the following networking issues:

- **Addressing routing issues**
  There are a variety of routing configuration issues that you need to address. If you did not create a default route with the First-Time Boot utility, you must configure a default route for the BIG-IP Controller. You also must set up routes for the nodes that the BIG-IP Controller manages. You may also want to configure GateD, which allows dynamic routing information to automatically be updated on the BIG-IP Controller.
- **Configuring DNS on the BIG-IP Controller**
  You may need to configure the BIG-IP Controller for DNS resolution or for DNS proxy, and you may even need to convert from rotary or round robin DNS.

- **Configuring email on the BIG-IP Controller**
  There are some special requirements that you need to take into account when configuring email on the BIG-IP Controller.

Addressing routing issues

The BIG-IP Controller must communicate properly with network routers, as well as the servers, firewalls, and other routers that it manages. Because there is a variety of router configurations, and varying levels of direct control an administrator has over each router, you need to carefully review the router configurations in your own network. You may need to change some routing configurations before you put the BIG-IP Controller into production.

The BIG-IP Controller supports static route configurations, dynamic routing (via BGP4, RIP1, RIP2, and OSPF), and subnetting. However, the BIG-IP Controller is also designed to eliminate the need for you to modify routing tables on a router that routes to a BIG-IP Controller. Instead, the BIG-IP Controller uses Address Resolution Protocol (ARP) to notify routers of the IP addresses that it uses on each interface, as well as on its virtual servers.

The following sections address these common routing issues:

- Routing from a BIG-IP Controller to a gateway to the external network
- Routing from content servers to the BIG-IP Controller
- Routing from a BIG-IP Controller to content servers that are on different logical networks
- Setting up dynamic routing with GateD

Routing from a BIG-IP Controller to a gateway to the external net

The BIG-IP Controller needs a route to the external network. For most configurations, this should be configured as the **default** route on the BIG-IP Controller.

◆ Note

*For multiple gateways to the external network, you can configure a last hop pool. For more information, see the **Reference Guide**.*

During installation, you were prompted to configure a default route for the BIG-IP Controller. If you need to change the default route at this time, you can set a new default route by editing the **/etc/netstart** file.

To change the default route

1. Open the **/etc/netstart** file in a text editor, such as **vi** or **pico**.
2. Change the default route entry using the following syntax:

   ```
   defroute="<router IP>"
   ```

3. Save and close the file.
4. Reboot the BIG-IP Controller.

Routing from content servers to the BIG-IP Controller

The content servers being load balanced by the BIG-IP Controller need to have a default route set to the internal IP alias (source processing) of the BIG-IP Controller. For most configurations, this should be configured as the **default** route on the content server.

For information about setting the default route for your content servers, refer to the product documentation for your server.

Routing between a BIG-IP Controller and content servers on different logical networks

If you need to configure the BIG-IP Controller to use one or more nodes that actually sit on a different logical network from the BIG-IP Controller, you need to assign one or more additional routes to get to those nodes. Set each node's default route in such a way that traffic goes back through the BIG-IP Controller internal interface.

In the following examples, the nodes are on 192.168.6/24 and the BIG-IP Controller internal interface is on 192.168.5/24. There are two possible situations which you may have to address:

- 192.168.5/24 and 192.168.6/24 are on the same LAN (either sharing media or with a switch or hub between them).
- 192.168.5/24 and 192.168.6/24 are on two different LANs with a router between them.

Case 1: Same LAN

If the nodes are on the same LAN as the BIG-IP Controller, you simply need to add an interface route for 192.168.6/24 to the BIG-IP Controller's internal interface. You can add this route to the bottom of the **/etc/rc.local** file using the following syntax:

```
route add -net 192.168.6 -interface exp1
```

◆ Note

*You must have the interface defined correctly in the **/etc/hosts** file in order to use this syntax.*

Case 2: Different LANs

If you have nodes on different LANs from the BIG-IP Controller, you need to add a static gateway route on the BIG-IP Controller itself. For example:

```
route add -net 192.168.6.0 -gateway 192.168.5.254
```

You also need to set the default route on the nodes to point to the router between the LANs. For example:

```
route add default -gateway 192.168.6.254
```

Finally, you need to set the default route on the router between the LANs to the BIG-IP Controller's shared alias. For example, type the command:

```
route add default -gateway 192.168.5.200
```

◆ Note

These examples assume you are using a UNIX-based router. The exact syntax for your router may be different.

It is not necessary to set the default route for nodes directly to the BIG-IP Controller, as long as the default path eventually routes through the BIG-IP Controller.

Setting up dynamic routing with GateD

The GateD daemon allows the BIG-IP Controller to exchange dynamic routing updates with your routers. Setting up the GateD daemon is a three-part task:

- ❖ You need to create the GateD configuration file, **/etc/gated.conf**.
- ❖ You need to start the GateD daemon.
- ❖ You need to edit the **/etc/netstart** file.

◆ Tip

You are not required to configure GateD on the BIG-IP Controller. The BIG-IP Controller can meet most routing requirements without using GateD.

◆ Note

Additional documentation for GateD is available through the web server on the BIG-IP Controller.

To create the GateD configuration file

GateD relies on a configuration file, typically named **/etc/gated.conf**, which can be relatively simple, or can be very complex, depending on the routing needs of your network. The BIG-IP web server includes the GateD online documentation (in the

Configuration utility home page, under *Online Documentation* section, click **GateD**). Note that the GateD configuration guide details the process of creating the GateD configuration file, and also provides samples of common protocol configurations.

To immediately start the GateD daemon on the BIG-IP Controller

Once you create the GateD configuration file, you need to start the GateD daemon on the command line using the following command:

```
bigip# gated
```

To enable starting GateD each time the BIG-IP Controller starts

To start GateD each time the BIG-IP Controller starts, change the **gated** variable in the **/etc/netstart** file as shown below:

```
gated=YES
```

Configuring DNS on the BIG-IP Controller

If you plan to use DNS in your network, you can configure DNS on the BIG-IP Controller. There are three different DNS issues that you may need to address when setting up the BIG-IP Controller:

- Configuring DNS resolution on the BIG-IP Controller
- Configuring DNS proxy
- Converting from rotary or round robin DNS

Configuring DNS resolution

When entering virtual addresses, node addresses, or any other addresses on the BIG-IP Controller, you can use the address, host name, or fully qualified domain name (FQDN).

The BIG-IP Controller looks up host names and FQDNs in the **/etc/hosts** file. If it does not find an entry in that file, then it uses DNS to look up the address. In order for this to work, you need to create an **/etc/resolv.conf** file. The file should have the following format:

nameserver <DNS_SERVER_1>

search <DOMAIN_NAME_1> <DOMAIN_NAME_2>

In place of the **<DNS_SERVER_1>** parameter, use the IP address of a properly configured name server that has access to the Internet. You can specify additional name servers as backups by inserting an additional **nameserver** line for each backup name server.

If you configure the BIG-IP Controller itself as a DNS proxy server, then we suggest that you choose its loopback address (**127.0.0.1**) as the first name server in the **/etc/resolv.conf** file.

Replace the **<DOMAIN_NAME_1>** and **<DOMAIN_NAME_2>** parameters with a list of domain names to use as defaults. The DNS uses this list to resolve hosts when the connection uses only a host name, and not an FQDN. When you enter domain names in this file, separate each domain name with a space, as shown.

A sample configuration file is shown in Figure 3.2.

```
; example /etc/resolv.conf
nameserver 127.0.0.1
nameserver 127.16.112.2  ;ip address of main DNS server
search mysite.com store.mysite.com
```

Figure 3.2 *Sample **/etc/resolv.conf** file*

You can also configure the order in which name resolution checks are made by configuring the **/etc/irs.conf** file. You should set this file so that it checks the **/etc/hosts** file first, and then checks for DNS entries. See Figure 3.3, for an example of how to make the entry in the **/etc/irs.conf** file.

```
hosts           local   continue
hosts           dns
```

Figure 3.3 *Sample entry for the **/etc/irs.conf** file*

Configuring DNS proxy

The BIG-IP Controller is automatically configured as a DNS proxy or forwarder. This is useful for providing DNS resolution for servers and other equipment load balanced by the BIG-IP Controller.

To re-configure DNS proxy, you simply edit the **/etc/named.boot** file that contains these two lines:

```
forwarders <DNS_SERVERS>
options forward-only
```

In place of the **<DNS_SERVER>** parameter, use the IP addresses of one or more properly configured name servers that have access to the Internet.

You can also configure the BIG-IP Controller to be an authoritative name server for one or more domains. This is useful when DNS is needed in conjunction with internal domain names and network addresses for the servers and other equipment behind the BIG-IP Controller. Refer to the BIND documentation for more details.

Converting from rotary or round robin DNS

If your network is currently configured to use rotary DNS, your node configuration may not need modification. However, you need to modify your DNS zone tables to map to a single IP address instead of to multiple IP addresses.

For example, if you had two Web sites with domain names of **www.SiteOne.com** and **www.SiteTwo.com**, and used rotary DNS to cycle between two servers for each Web site, your zone table might look like the one in Figure 3.4.

```
www.SiteOne.com  IN A 192.168.1.1
                 IN A 192.168.1.2
www.SiteTwo.com  IN A 192.168.1.3
                 IN A 192.168.1.4
```

Figure 3.4 *Sample zone table with two Web sites and four servers*

In the BIG-IP Controller configuration, the IP address of each individual node used in the original zone table becomes hidden from the Internet. We recommend that you use the Internet reserved address range as specified by RFC 1918 for your nodes. In place of multiple addresses, simply use a single virtual server associated with your site's domain name.

Using the above example, the DNS zone table might look like the zone table shown in Figure 3.5.

```
www.SiteOne.com  IN A 192.168.100.231
www.SiteTwo.com  IN A 192.168.100.232
```

Figure 3.5 *Sample zone table with two Web sites and two servers.*

Configuring Email

Another optional feature you can set up when you configure the BIG-IP Controller is email. You can configure the BIG-IP Controller to send email notifications to you, or to other administrators. The BIG-IP Controller uses Sendmail as its mail transfer agent. The BIG-IP Controller includes a sample Sendmail configuration file that you can use to start with, but you will have to customize the Sendmail setup for your network environment before you can use it.

Before you begin setting up Sendmail, you may need to look up the name of the mail exchanger for your domain. If you already know the name of the mail exchanger, skip to ***Setting up Sendmail*** on page 3-16.

Finding the mail exchanger for your domain

You can use the **nslookup** command on the BIG-IP Controller or any workstation that is configured for DNS **lookup**. Once you find the primary IP address for your domain, you can find the mail exchanger for your domain.

To find the mail exchanger

1. First you need to identify the default server name for your domain. From a workstation capable of name resolution, type the following on the command line:

   nslookup

2. The command returns a default server name and corresponding IP address:
   Default Server: <server name>
   Address: <server addr>

3. Next, use the domain name to query for the mail exchanger:

   set q=mx
   <domain name>

The information returned includes the name of the mail exchanger. For example, the sample information shown in Figure 3.6 lists **bigip.net** as the preferred mail exchanger.

```
bigip.net   preference = 10, mail exchanger = mail.SiteOne.com
bigip.net   nameserver = ns1.bigip.net
bigip.net   nameserver = ns2.bigip.net
bigip.net   internet address = 192.17.112.1
ns1.bigip.net       internet address = 192.17.112.2
ns2.bigip.net       internet address = 192.17.112.3
```

Figure 3.6 *Sample mail exchanger information*

Setting up Sendmail

When you actually set up Sendmail, you need to open and edit a couple of configuration files. Note that the BIG-IP Controller does not accept email messages, and that you can use the **crontab** utility to purge unsent or returned messages, and that you can send those messages to yourself or another administrator.

To set up and start Sendmail

1. Copy **/etc/sendmail.cf.off** to **/etc/sendmail.cf**.

2. To set the name of your mail exchange server, open the **/etc/sendmail.cf** and set the DS variable to the name of your mail exchanger. The syntax for this entry is:

   ```
   DS<MAILHUB_OR_RELAY>
   ```

3. Save and close the **/etc/sendmail.cf** file.

4. To allow Sendmail to flush outgoing messages from the queue for mail that cannot be delivered immediately, open the **/etc/crontab** file, and change the last line of the file to read:

```
0,15,30,45 * * * *   root /usr/sbin/sendmail -q > /dev/null 2>&1
```

5. Save and close the **/etc/crontab** file.

6. To prevent returned or undelivered email from going unnoticed, open the **/etc/aliases** file and create an entry for **root** to point to you or another administrator at your site.

   ```
   root: networkadmin@SiteOne.com
   ```

7. Save and close the **/etc/aliases** file.

8. You now need to run the **newaliases** command to generate a new aliases database that incorporates the information you added to the **/etc/aliases** file.

9. To turn Sendmail on, either reboot the system or type the following command:

   ```
   /usr/sbin/sendmail -bd -q30m
   ```

Using a serial terminal with the BIG-IP Controller

There are several different ways to add a serial terminal to the BIG-IP Controller. You can add a serial terminal in addition to the console, or you can add a serial terminal as the console. The difference between the two is:

- A serial terminal configured as a terminal displays a simple login. You can log in and run commands and edit files. In this case, you can use the serial terminal in addition to the keyboard and monitor.

- A serial terminal configured as the console displays system messages and warnings in addition to providing a login prompt. In this case, the serial terminal replaces the keyboard and monitor.

Connect a serial line cable between the terminal device and the BIG-IP Controller. Note: On the back of BIG-IP is a male, 9-Pin RS232C connector labeled "Terminal". (Be sure not to confuse this with the failover connection which is also a male, 9-pin connector.) The connector is wired as a DTE device, and uses the signals described in Table 3.1.

Pin	Source	Usage
1	External	Carrier detect
2	External	Received data
3	Internal	Transmitted data
4	Internal	Data terminal ready
5	Both	Signal ground
7	Internal	Request to send
8	External	Clear to send

Table 3.1 *Serial line cable signals*

The connector is wired for direct connection to a modem, with receipt of a Carrier Detect signal generating transmission of a login prompt by BIG-IP. If you are planning to connect to a terminal or to connect a PC and utilize a terminal emulation program such as HyperTerminal™, you will need a null modem cable with the wiring to generate the signals shown in Table 3.1.

◆ Note

You can achieve acceptable operation by wiring pins 7 to 8 and pins 1 to 4 at the back of BIG-IP Controller (and turning hardware flow control ***off*** *in your terminal or terminal emulator).*

To configure a serial terminal in addition to the console

If you want to configure a serial terminal for the BIG-IP Controller in addition to the standard console, use the following procedure.

1. Configure the serial terminal settings in your terminal or terminal emulator or modem as follows:
   - 9600 baud
   - 8 bits
   - 1 stop bit
   - No parity
2. Open the **/etc/ttys** file and find the line that reads **tty00 off**. Modify it as shown here:

```
# PC COM ports (tty00 is DOS COM1)
tty00 "/usr/libexec/getty default" vt100 in secure
```

3. Save the **/etc/ttys** file and close it.
4. Reboot the BIG-IP Controller.

To configure a serial terminal as the console

In order to configure the serial terminal as the console, follow these steps:

1. Disconnect the keyboard from the BIG-IP Controller
2. Connect the serial terminal to the BIG-IP Controller. When there is no keyboard connected to the BIG-IP Controller, the BIG-IP Controller defaults to using the serial port for the console.
3. Reboot the controller

To force a serial terminal to be the console

In the case where you have not connected the serial terminal or it is not active when the BIG-IP Controller is booted, as it might be if you are using a terminal server or dial-up modem, you can force the controller to use the serial terminal as a console. To do this, follow this procedure:

1. Edit the **/etc/boot.default** file.
   Find the entry **-console auto**. Change this entry to **-console com**.
2. Save the **/etc/boot.default** file and exit the editor.
3. Plug the serial terminal into the serial port on the BIG-IP Controller.
4. Turn on the serial terminal.
5. Reboot the controller.

WARNING

You can only access the BIG-IP Controller through the serial port or the network if you perform this procedure.

1) Once the serial line has been configured as the console, keyboard/monitor access is disabled. Once this occurs, logins are only possible via Secure Telnet (SSH), if configured, or the serial line.

2) Be careful when editing boot.default. If this file is corrupted, the system will not boot at all. Save a backup copy of the original file and be prepared with a bootable CD-ROM.

3) Be sure that boot.default contains either the line: "-console com" or the line: "-console auto".

Note

You do not need to disconnect the keyboard if you use this procedure to force the serial line to be the console.

Configuring RADIUS authentication

You can configure the BIG-IP Controller to use a RADIUS server on your network to authenticate users attempting to access the controller with SSH. This allows you to use the RADIUS server as a central repository of users that can access the BIG-IP Controller for administrative purposes.

To do this, configure the BIG-IP Controller to act as a Network Access Server (NAS) for a RADIUS server in your network. When you set up this feature, client connections received by the BIG-IP Controller for users not listed in the local account database are routed to the RADIUS server to be authenticated. If the user is authenticated, the user is logged in as the BIG-IP Controller user that you specify in the RADIUS user setting.

◆ Note

RADIUS authentication through the BIG-IP Controller is based on the username/password only. Challenge-response authentication methods are not supported.

You can configure the BIG-IP Controller to use either version 1.3.7 or version 2.0.12.1, or both, of the **sshd** for SSH authentication.

◆ Note

*If you want to support only SSH version 1.x clients, configure **sshd** version 1.3.7. Do not configure **sshd** version 2.0.12.1. However, if you want to support version 1.x and version 2.x clients, configure **sshd** version 2.0.12.1.*

RADIUS ports on the BIG-IP Controller

The BIG-IP Controller uses the ports **1645/udp** for communicating with the RADIUS server. If your RADIUS server uses different ports, such as **1812/udp**, you must change the ports used by the BIG-IP Controller to these ports. To do this, use a text editor such as **vi** or **pico** to change the existing RADIUS port entry in the **/etc/services** file on each BIG-IP Controller:

```
radius          1812/tcp                        # Radius
radacct         1813/udp                        # Radius Accounting
```

Figure 3.7 *Alternative ports on the BIG-IP Controller for the RADIUS server*

Configuring sshd version 1.3.7

You can configure version 1.3.7 of the **sshd** by editing the **/etc/sshd_config** on the BIG-IP Controller with **pico** or **vi**. The following entries must be in the **sshd_config** file:

- **RadiusServer**
  This entry is the host name or IP address of the RADIUS server.
- **RadiusKey**
  This entry is the shared secret key of the RADIUS server. This key should be at least 16 characters long.
- **RadiusNasIP**
  This is the host name or IP address of the interface on the BIG-IP Controller connected to the network that hosts the RADIUS server. Note that you can only use interfaces set to **admin port open** for RADIUS authentication.

❖ **RadiusUser**
This entry is the user name of the local BIG-IP Controller user, such as root. When the RADIUS user is authenticated, the user is logged into the controller as this user.

Note

*The most secure method for using RADIUS with the BIG-IP Controller is to create a **RadiusUser** entry that has a low level of privileges. After you are authenticated and you log in to the BIG-IP Controller as the low privilege user, use the **su** command to gain root privileges.*

WARNING

*For security reasons, we recommend that you use IP addresses instead of host names for the entries in this file. If you specify a host name for an entry, we recommend that you add the host name to the **/etc/hosts** file.*

For example, Figure 3.8 is an example of the entries you might make in the **sshd_config** file on the BIG-IP Controller:

```
RadiusServer 12.34.56.78
RadiusKey    my_radius_server.key
RadiusNasIP  172.16.42.200
RadiusUser   radius_user
```

Figure 3.8 *Example entries from the **sshd_config** file*

Configuring sshd version 2.0.12.1

You can configure version 2.0.12.1 of the **sshd** by editing the **/etc/ssh2/sshd2_config** on the BIG-IP Controller with **pico** or **vi**. The following entries must be in the **sshd2_config** file:

❖ **RadiusServer**
This entry is the host name or IP address of the RADIUS server.

- **RadiusKey**
  This entry is the shared secret key of the RADIUS server. This key should be at least 16 characters long.
- **RadiusNasIP**
  This is the host name or IP address of the interface on the BIG-IP Controller connected to the network that hosts the RADIUS server. Note that you can only use interfaces set to **admin port open** for RADIUS authentication.
- **RadiusUser**
  This entry is the user name of the local BIG-IP Controller user, such as root. When the RADIUS user is authenticated, the user is logged into the controller as this user.

◆ Note

*The most secure method for using RADIUS with the BIG-IP Controller is to create a **RadiusUser** entry that has a low level of privileges. After you are authenticated and you log in to the BIG-IP Controller as the low privilege user, use the **su** command to gain root privileges.*

To support SSH version 1.x clients, you must add the following entries to the **/etc/ssh2/sshd2_config** file.

- **Ssh1Compatibility**
  This parameter must be set to **yes**.
- **Sshd1Path**
  This entry is the path to **sshd** version 1. In this case, the path is **/usr/local/sbin/sshd1**.

For example, Figure 3.9 is an example of the entries you might make in the **sshd2_config** file on the BIG-IP Controller:

```
RadiusServer 12.34.56.78
RadiusKey    my_radius_server.key
RadiusNasIP  172.16.42.200
RadiusUser   radius_user

Sshd1Compatibility yes
Sshd1Path /usr/local/bin/sshd1
```

Figure 3.9 *Example entries from the* ***sshd2_config*** *file*

F5 Networks

BIG·IP®

Local Traffic Controller

Index

/etc/bigip.conf file 2-13
/etc/ethers file 2-13
/etc/hosts file 2-13, 3-1, 3-2, 3-3, 3-13
/etc/irs.conf file 3-13
/etc/netstart file 2-13
/etc/resolv.conf file 3-13
/etc/sshd_config file 2-13
/var/f5/httpd/conf/httpd.conf file 2-13

A

Address Resolution Protocol (ARP). See ARP
ARP 3-8
authoritative name servers 3-14

B

BIG/IP Controller types I-7
BIG/pipe utility I-2
BIG/top utility
    described I-2
BIG-IP Controller types 1-3

C

cache server types I-6
cacheable content determination
    defined I-6
command line access 3-3–3-7
config_ftpd script 3-4
configuration settings 2-12–2-13
Configuration utility
    described I-1
Configuration utility requirements I-2
connections
    administrative 2-10
content retrieval I-6
content servers 3-9
content types
    for caching I-6

D

Data Fellows 1-1, 3-3
default route configuration 2-4, 2-6
DNS configuration 3-12–3-15
    configuring proxy 3-14
    converting from rotary 3-14
    converting from round robin 3-14
    proxy forwarding 2-4
    resolving names 3-12–3-13
    zone tables 3-14–3-15
domain names 3-14

E

Earth ground 1-7
email configuration 3-15–3-17
encrypted connections I-8
external interfaces 2-5–2-8
external network interfaces. See external interfaces

F

fail-over cable 1-2, 1-9
First-Time Boot utility 2-1–2-13
    defined I-1
    required information 2-1
FQDN 3-12
F-Secure SSH client option 3-3–3-7
    and encrypted communications I-8
    as a remote shell I-2
    documentation 1-1, 3-4
    downloading 3-4
    downloading with FTP 3-4

installing on UNIX 3-7
installing on Windows 95 or NT 3-6
transferring 3-5

G

GateD 3-11

H

hardware
usage guidelines 1-6

hardware installation
connecting 1-8
of 2U controller 1-5–1-6
of 4U controller 1-3–1-4
planning 1-6
procedures 1-8

hardware requirements
components 1-1
peripherals 1-2

host names
defining 2-3
defining additional 3-1

HTTP request headers
and content caching I-6

httpd.conf file 2-12

I

intelligent cache population
defined I-6

interface cards. See NICs

interface configuration 2-5–2-9
attributes 2-5

interfaces
external. See external interfaces
internal. See internal interfaces

internal interfaces 2-5–2-6, 2-8–2-9
choosing primary 2-9

internal network interfaces. See internal interfaces

internal network. See internal interfaces

IP addresses
changing 2-11
configuring default route 2-4
configuring fail-over 2-1, 2-5
defining I-1
external interfaces. See external interfaces.
internal interfaces. See internal interfaces

IP aliases 2-6

L

LED indicators 2-7

license statement 2-2

lithium battery 1-7

load balancing
configuring I-1
monitoring I-1

M

mail exchanger 3-16

media types 2-9

MIB I-2

monitoring methods
and command-line utilities I-2

N

name servers 3-14

netmask 2-6

network adapters 2-7

network addresses 3-14

network interface cards (NICs). See NICs

NICs
installing 1-9

non-transparent cache servers
described I-6
nslookup command 3-15

P

ports 1-3–1-6
power cable 1-9
procedures
changing default router configuration 3-9
downloading F-Secure SSH client 3-4–3-6
installing hardware 1-8
setting up F-Secure SSH client 3-6–3-7
setting up sendmail 3-16

R

rack mounting 1-6
RADIUS authentication 3-22, 3-23, 3-25
reconfig-httpd utility 2-11
redundant systems
choosing fail-over IP addresses 2-5
configuring external interfaces 2-7
configuring internal interfaces 2-9
hardware-based 1-2
selecting unit ID 2-5
remote administration 1-2
root password
defining I-1
root password definition 2-3
rotary DNS. See DNS configuration
round robin 3-14
See also DNS configuration
router configurations 3-8–3-12
default 3-9
examples 3-10
from content servers 3-9
on different logical networks 3-10
with GateD 3-11
routers 3-1

S

secure connections I-8
security
encryption. See encryption
for web server 2-11
sendmail 3-16
SNMP MIB I-2
SSH client option
See F-Secure SSH client option
SSH client. See F-Secure SSH client option
sshd version 1.3.7 3-22
sshd version 2.0.12.1 3-23
SSL connections I-8
system setup I-1

T

time zone configuration 2-4

U

US functionality
generating authentication certificates 2-11
support for encryption 1-1
utilities I-2
First-Time Boot 2-1–2-13
reconfig-httpd 2-11

V

ventilation 1-6
virtual servers
adding host names 3-1
See also IP addresses

voltage 1-7

W

web server access I-1
web server settings 2-11
workstation configuration 3-3–3-7