f5 NETWORKS
BIGip
Getting Started

BIG/ip™ Controller
Getting Started Guide

version 3.0

Service and Support Information

Product Version

This manual applies to version 3.0 of the BIG/ip® Controller platform.

Obtaining Technical Support

Web	tech.f5.com
Phone	(206) 505-0888
Fax	(206) 505-0802
Email (support issues)	support@f5.com
Email (suggestions)	feedback@f5.com

Contacting F5 Networks

Web	www.f5.com
Toll-free phone	(888) 88BIG-IP
Corporate phone	(206) 505-0800
Fax	(206) 505-0801
Email	sales@f5.com
Mailing Address	200 1st Avenue West Suite 500 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.

Copyright© 2000 by
F5 Networks, Inc.
Seattle, Washington
All rights reserved. Printed in U.S.A.

Trademarks

F5, BIG/ip, and 3DNS are registered trademarks of F5 Networks, Inc. see/IT and global/SITE 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® Controller, if originally sold within the United States, may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this BIG/ip® 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.

Canada Compliance

This class A digital apparatus complies with Canadian ICES-003.

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.

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/).

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 (for example, BIG/ip V1.5 to BIG/ipV1.6). 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 (for example, BIG/ip V1.x to BIG/ip V2.0). 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 for the BIG/ip® Controller, whether such software is provided separately or as an integral part of a hardware product. As used herein, the term "Software" will refer to all such software, and the corrections, updates, new releases and new versions of such software. A product that consists of Software only will be referred to as a "Software Product" and a combination Software/Hardware product will be referred to as a "Combination Product." All Software is licensed, not sold, by F5. 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 solely on a single central processing unit owned or leased by Licensee. 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. Licensee may make one back-up copy of any Software Product, provided the back-up copy contains the same copyright and proprietary information notices as the original Software Product. Licensee is not authorized to copy the Software contained in

a Combination Product. The Software incorporates certain third party software which is used subject to licenses from the respective owners.

3. Restrictions. The Software, documentation, and the associated copyrights are owned by F5 or its licensors, and are protected by law and international treaties. Except as provided above, 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 laws and governmental 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 exporting or using cryptographic software.

5. Limited Warranty.

   a) 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.

   b) 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 F5 product that fails during the warranty

period at no cost to Licensee. Any 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 Licensee, at F5's expense, no later than 7 days after receipt by F5. Title to any returned product 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.

c) Restrictions. The foregoing limited warranties extend only to the original Licensee, and do not apply if a Software Product or Combination Product (i) has been altered, except by F5, (ii) has not been installed, operated, repaired, or maintained in accordance with F5's instructions, (iii) has been subjected to abnormal physical or electrical stress, misuse, negligence, or accident or (iv) 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. 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 THE SOFTWARE, 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 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. 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 THE SOFTWARE OR OTHER GOODS OR SERVICES FURNISHED TO LICENSEE BY F5, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. Termination. This License is effective until terminated, and will automatically terminate if Licensee fails to comply with any of its provisions. Upon termination of this License, the Licensee will destroy the Software and documentation and all copies or portions thereof.

8. 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

Chapter 1

Introduction to the BIG/ip Controller Getting Started Guide

Welcome to the BIG/ip Controller Getting Started Guide 1-1
BIG/ip Controller specifications 1-1
Configuration scalability 1-2
BIG/ip Controller platform options 1-2
Finding help and technical support resources 1-3

Chapter 2

Setting up the Hardware

Unpacking and installing the hardware 2-1
Reviewing the hardware requirements 2-1
Familiarizing yourself with the BIG/ip Controller hardware 2-3
Environmental requirements and usage guidelines 2-7
Installing and connecting the hardware 2-8
Running the First-Time Boot utility 2-11
Gathering the information 2-12
Starting the First-Time Boot utility 2-12
Defining a root password 2-13
Defining a host name 2-14
Configuring a default route 2-14
Configuring a time zone 2-14
Configuring the DNS forwarding proxy settings 2-14
Configuring the interfaces 2-14
Configuring remote administration 2-19
Configuring settings for the BIG/ip web server 2-20
Confirming your configuration settings 2-22
Committing your configuration settings to the system 2-22
Defining additional host names 2-23
Preparing workstations for command line access 2-25
Downloading the F-Secure SSH client from the BIG/ip web server . . 2-26
Downloading the F-Secure SSH client using FTP 2-26
Setting up the F-Secure SSH client on a Windows 95 or Windows NT workstation 2-28
Setting up the F-Secure SSH client on a UNIX workstation 2-29

Chapter 3

Getting Started with a Basic Configuration

Setting up a basic configuration 3-1

Configuring virtual servers 3-4
Using standard or wildcard virtual servers 3-4
Using additional features with virtual servers 3-5
Defining standard virtual servers 3-6
Defining wildcard virtual servers 3-8
Allowing access to ports and services 3-12
Configuring the timer settings 3-13
Setting the node ping timer 3-14
Setting the timer for reaping idle connections 3-15
Setting the service check timer 3-17
Service checking for wildcard servers and ports 3-18
Changing the global load balancing mode 3-19
Using Ratio mode 3-20
Configuring NATs and IP forwarding for nodes 3-22
Defining a standard network address translation (NAT) 3-24
Defining a secure network address translation (SNAT) 3-25
Setting up IP forwarding 3-28
Configuring Extended Content Verification service checking 3-29
ECV service check properties 3-30
Writing regular expressions for ECV service checks 3-31
Setting up ECV service check in the F5 Configuration utility 3-33
Manually configuring and testing the /etc/bigd.conf file 3-34
Configuring persistence for e-commerce and other dynamic content sites . 3-36
Setting up SSL persistence 3-38
Setting up simple persistence 3-39
Configuring and synchronizing redundant systems 3-41
Preparing to use the synchronization command 3-42
Synchronizing configurations between controllers 3-43
Configuring fail-safe settings 3-44
Addressing general networking issues 3-46
Addressing routing issues 3-47
Configuring DNS on the BIG/ip Controller 3-51
Configuring Email 3-54
Basic configuration examples 3-56
A basic web site and e-commerce configuration 3-57
A basic intranet configuration 3-60

Index

1

Introduction to the BIG/ip Controller Getting Started Guide

- **Welcome to the BIG/ip Controller Getting Started Guide**
- **BIG/ip Controller specifications**
- **Finding help and technical support resources**

Welcome to the BIG/ip Controller Getting Started Guide

Welcome to the *BIG/ip® Controller Getting Started Guide*. This guide describes how to set up the BIG/ip Controller hardware and how to use the First-Time Boot utility for basic software configuration. This book is the first in a series of three guides:

- ***BIG/ip Controller Getting Started Guide***
  Use this guide for hardware configuration and basic software configuration.
- ***BIG/ip Controller Administrator Guide***
  Use this guide for advanced software configuration and administration of the BIG/ip Controller.
- ***BIG/ip Controller Reference Guide***
  Use this guide for reference information including the BIG/pipe command line commands, BIG/ip configuration utilities, and system utilities.

BIG/ip Controller specifications

The BIG/ip Controller is a network appliance that manages and balances traffic for networking equipment such as web servers, cache servers, routers, firewalls, and proxy servers. A variety of useful features meets the special needs of e-commerce sites, Internet service providers, and managers of large intranets. The system is highly configurable, and its web-based and command line configuration utilities allow for easy system set up and monitoring.

Adding a BIG/ip Controller to your network ensures that your network remains reliable. The BIG/ip Controller continually monitors the servers and other equipment it manages, and never attempts to send connections to servers that are down or too busy to handle the connection. The BIG/ip Controller uses a variety of methods to monitor equipment, from simple pings to more advanced methods, such as Extended Content Verification that

verifies whether a server returns specific site content. The BIG/ip Controller also offers several layers of redundancy that ensure its own reliability.

Configuration scalability

The BIG/ip Controller is a highly scalable and versatile local area load balancing solution. You can actually configure a single BIG/ip Controller to manage up to 10,000 virtual servers, though most common configurations are significantly smaller. The number of servers, firewalls, or routers that a single BIG/ip Controller can load balance is limited only by the capacity of your network media, such as Ethernet. The BIG/ip Controller supports a variety of media options, including Fast Ethernet, Gigabit Ethernet, and FDDI. The maximum number of concurrent connections that a BIG/ip Controller can manage is determined by the amount of RAM in your particular BIG/ip Controller hardware configuration.

BIG/ip Controller platform options

The BIG/ip Controller platform offers three different systems, each of which can run as a stand-alone controller, or as a redundant controller pair:

- **The BIG/ip LB Controller**
  The BIG/ip LB Controller provides basic load balancing features. Note that the BIG/ip LB Controller does not support all of the features documented in this guide. BIG/ip LB Controllers distributed in the US also support encrypted administrative connections using SSL for connections to the web-based F5 Configuration utility.
- **The BIG/ip HA Controller**
  In addition to the basic load balancing features supported on the BIG/ip LB Controller, the BIG/ip HA Controller supports advanced features, such as Extended Content Verification, and also supports high-end security for administrative shell connections. BIG/ip HA Controllers distributed in the US also

support encrypted administrative connections using SSH for shell connections and SSL for connections to the web-based F5 Configuration utility.

- **The BIG/ip HA+ Controller**
  The BIG/ip HA+ Controller supports the same features as the BIG/ip HA Controller. In addition to the software features, the HA+ controller offers high-end hardware for high traffic sites.

◆ Note

BIG/ip Controllers distributed outside of the United States, 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.

Finding help and technical support resources

In addition to this getting started guide, you can find technical documentation about the BIG/ip Controller in the following locations:

- **BIG/ip Controller Administrator Guide**
  Use this guide for advanced software configuration and administration of the BIG/ip Controller.
- **BIG/ip Controller Reference Guide**
  Use this guide for reference information including the BIG/pipe command line commands, BIG/ip configuration utilities, and system utilities.
- **Release notes**
  The release note for the current version of the BIG/ip Controller is available on the BIG/ip web server. 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. You can obtain the latest version of the release notes from http://tech.f5.com.

- **Online help for BIG/ip Controller features**
  You can find help online in three different locations:
  - The BIG/ip web server has a PDF version of this guide. Note that some BIG/ip Controller upgrades replace the online administrator guide with an updated version of the guide.
  - The web-based F5 Configuration utility has online help for each screen. Simply click the **Help** button on the toolbar.
  - Individual BIG/pipe commands have online help, including command syntax and examples, in standard UNIX man page format. Simply type the command key word 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 BIG/ip web server 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, and updates for administrator guides (in PDF format). To access this site, you need to obtain a customer ID and a password from the F5 Help Desk.

2

Setting up the Hardware

- **Unpacking and installing the hardware**
- **Running the First-Time Boot utility**
- **Defining additional host names**
- **Preparing workstations for command line access**

Unpacking and installing the hardware

There are two basic tasks you must complete to get the BIG/ip Controller installed and set up.

- Connect the peripheral hardware and connect the BIG/ip Controller to the network.
- Turn the system on and run the First-Time Boot utility.
  The First-Time Boot utility is a wizard that helps you configure basic system elements such as administrative passwords, IP addresses, and host names for both the root system and for the BIG/ip web server. Once you complete the First-Time Boot utility, you can continue the configuration process either from a remote administrative workstation, or directly from the console.

In addition to these two basic tasks, you can perform the following tasks:

- Define additional host names for virtual servers and other devices on the network.
- Prepare workstations for command line access to the BIG/ip Controller.

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 Controller Getting Started Guide***
- One ***BIG/ip Controller Administrator Guide***
- One ***BIG/ip Controller Reference Guide***

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). Additionally, if you purchased a US BIG/ip Controller that supports encryption, you receive the ***F-Secure SSH Client*** manual, published by Data Fellows.

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 *To configure a serial terminal in addition to the console*, on page 2-10, 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 two hardware configurations. The LB and HA versions of the BIG/ip Controller ship in the 4U hardware configuration. The HA+ version of the BIG/ip Controller ships in a 2U hardware configuration. 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.

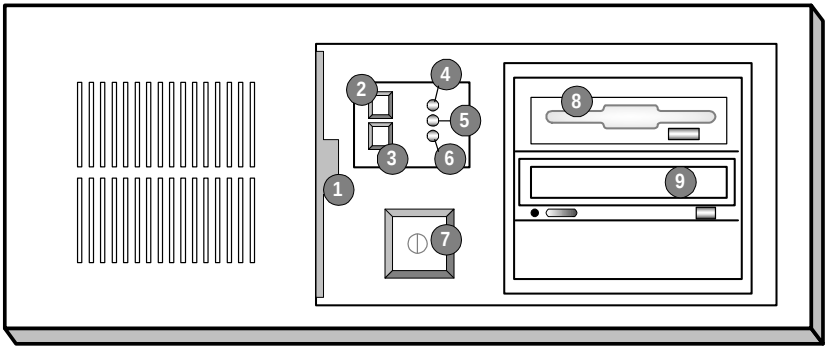

1. Fan filter	6. Power LED
2. Keyboard lock	7. On/off button
3. Reset button	8. 3.5 floppy disk drive
4. Keyboard lock LED	9. CD-ROM drive
5. Hard disk drive LED	

Figure 2.1 *Front view of a 4U BIG/ip Controller*

Figure 2.1 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 2.2, 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.

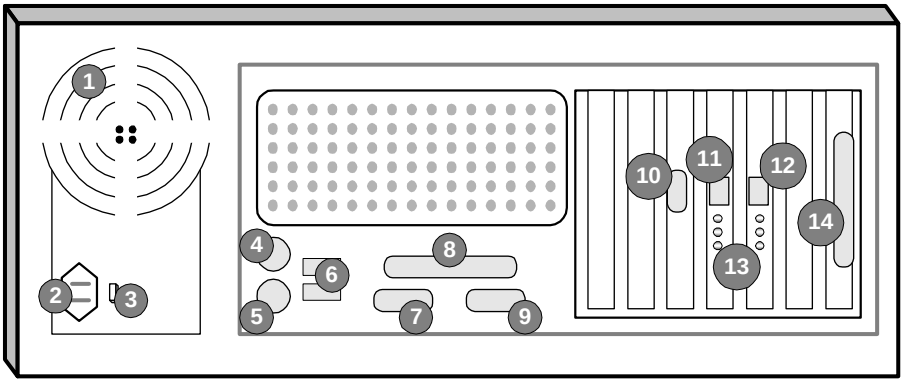

1. Fan	8. Printer port*
2. Power in	9. Fail-over port
3. Voltage selector	10. Video (VGA) port
4. Mouse port*	11. Internal interface (RJ-45)
5. Keyboard port	12. External interface (RJ-45)
6. Universal serial bus ports*	13. Interface indicator LEDs
7. Serial terminal port	14. Watchdog card*

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

Figure 2.2 *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.

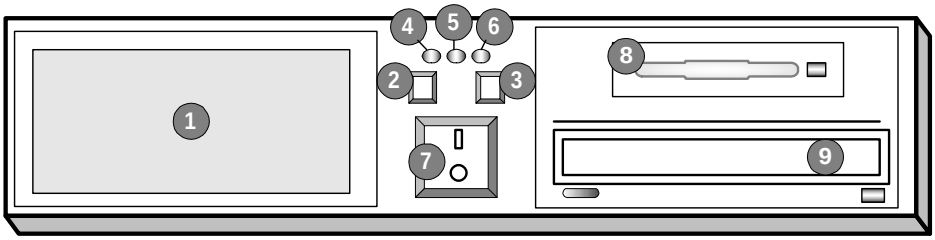

1. Fan filter	6. Power LED
2. Keyboard lock	7. On/off button
3. Reset button	8. Flash or PC card
4. Keyboard lock LED	9. CD-ROM drive
5. Hard disk drive LED	

Figure 2.3 *Front view of a 2U BIG/ip Controller*

Figure 2.3 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 2.4, 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.

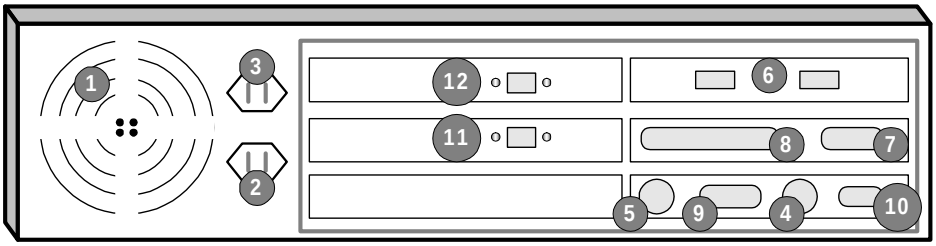

1. Fan
2. Power in
3. Power out
4. Mouse port*
5. Keyboard port
6. Universal serial bus ports*
7. Serial terminal port
8. Printer port*
9. Fail-over port
10. Video (VGA) port
11. Internal interface (RJ-45)
12. External interface (RJ-45)

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

Figure 2.4 *Back view of a 2U BIG/ip Controller*

Environmental requirements and usage 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, be sure to consider the following before you install the unit in the rack:

- You should always install the rack according to the manufacturer's instructions, and be sure to check the rack for stability before placing equipment in it.
- You should 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 50° C. The internal temperature of the room where the BIG/ip Controller is located should be considered for continued safe operation.

- Make sure that the branch circuit into which you plug the unit is not shared by more electronic equipment than it is designed to manage safely at one time.
- If you are installing the BIG/ip Controller in a location outside of the United States, you need to verify that the voltage selector is set appropriately before connecting the power cable to the unit.

WARNING

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

WARNING

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.

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 2.2 for 4U, or in Figure 2.4 for 2U) and the keyboard connector cable to the keyboard port (number 5 in Figure 2.2 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 2-1).

   - Optionally, if you are using a serial terminal as the console, connect the serial cable to the terminal serial port (number 7 in Figure 2.2 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 2.2 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 2.2 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 2.2 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 2.2 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,

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, you need to follow the configuration steps below. Note that if you are using a serial vt100 connection, you must edit the **/etc/ttys** file on the BIG/ip Controller.

Note

Before you configure the serial terminal, you must disconnect the keyboard from 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.

1. Configure the serial terminal settings 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
tty01 off
```

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

Running 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 set up.

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 BIG/ip web server

First, you configure all of the required information. Then you have the opportunity to confirm each individual setting or correct it if necessary. Then 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 root system and for the BIG/ip web server
- Host names for the root system 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 (as shown in Figures 2.1 and 2.3, number 7). 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, 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 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 (-), if you like. 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

Next, 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, you are also prompted 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, you are prompted 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 describes 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 ***BIG/ip Controller Administrator Guide****, Working with Special Features.*

If you have a redundant system, you are prompted 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.

You should 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 or 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**

You should 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

The screens that you see for configuring remote administration vary, depending on whether you have a US BIG/ip Controller, or an international BIG/ip Controller. On a US BIG/ip Controller, the first screen you see is the Configure SSH screen, which prompts you to type an IP address for SSH command line access. On international and BIG/ip LB Controllers that do not have SSH, the First-Time Boot utility skips this screen. Instead, you are prompted to configure access through Telnet and FTP.

When you configure shell access method, such as SSH, Telnet, or FTP, the First-Time Boot utility prompts you to create a support account for that method. You can use this support account to provide access to the BIG/ip Controller by an F5 Networks support engineer.

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. On US products, 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.

Note that 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 original 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.

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 2.5.

```
# 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 2.5 *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**). 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:

- BIG/ip Controllers distributed in the US support secure shell command line access using the F-Secure SSH client.
- BIG/ip Controllers distributed outside the US support command line access using a standard Telnet shell.

If you are working in the US with a 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, click the Downloads link. From the Downloads 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. Use this script to allow specific clients FTP or Telnet access to the BIG/ip Controller. However, this method is not recommended. 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. Start FTP:
   `ftp`
3. Open a connection to the remote workstation using 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. Switch to passive FTP mode:

   `passive`
6. Switch the transfer mode to binary:

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

8. Start the transfer process using 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.

3

Getting Started with a Basic Configuration

- **Setting up a basic configuration**
- **Configuring virtual servers**
- **Allowing access to ports and services**
- **Configuring the timer settings**
- **Changing the global load balancing mode**
- **Configuring NATs and IP forwarding for nodes**
- **Configuring Extended Content Verification service checking**
- **Configuring persistence for e-commerce and other dynamic content sites**
- **Configuring and synchronizing redundant systems**
- **Addressing general networking issues**
- **Basic configuration examples**

Setting up a basic configuration

This chapter covers the three essential configuration tasks that all users must complete, as well as the optional configuration tasks that most users find they want to do. Even if you want to use advanced features, such as IP filters or specialized load balancing modes, you start with the instructions in this chapter to set up your initial basic configuration.

A basic configuration sets up the BIG/ip Controller to do load balancing for one or more groups of content servers, firewalls, routers, or cache servers.

To set up a basic configuration, you need to do only the following three tasks.

- **Configure the virtual servers and the virtual server mappings**
  The virtual servers are the backbone of the system configuration, and they define the groups of servers or other network equipment that the BIG/ip Controller load balances.
- **Allow access to ports and services**
  The services and ports on a BIG/ip Controller are locked down and cannot accept connections until you specifically open them to network access. For each service that one or more of your virtual servers supports, you need to open the corresponding port number for network access. However, ports are automatically enabled when you use them in virtual server definition in the F5 Configuration utility.
- **Configure the timer settings**
  The BIG/ip Controller supports several timer settings, but for a simple configuration, there are only two that you need to set. First you need to set the amount of time that idle connections are allowed to remain open. Second, you need to set the frequency at which the BIG/ip Controller checks nodes to make sure they are up and available to accept connections passed on by a virtual server.

This chapter also covers additional configuration options that users typically add to a simple configuration, including:

- **Change the load balancing mode**
  You can use an alternate load balancing mode.

- **Configure NATs or IP forwarding**
  You can set up network address translation (NAT) or IP forwarding to allow direct connections to and from nodes.
- **Configure extended content verification service checking**
  You can configure extended content verification (ECV) to allow the BIG/ip Controller to verify that a server is responding to requests.
- **Set up persistence**
  You can set up persistence to accommodate e-commerce and other dynamic content sites that require returning clients to bypass load balancing and return to the same node to which they last connected.
- **Configure and synchronize redundant systems**
  You can configure and set up redundant BIG/ip Controller systems.

WARNING

When you set configuration options in the F5 Configuration utility, they are immediately saved to the appropriate configuration file. However, when you set configuration options using the BIG/pipe command line utility, they are temporarily stored in system memory, and are not saved to a configuration file unless you execute the ***bigpipe -s*** *command. For more information about this command, see the* ***BIG/ip Controller Reference Guide****, BIG/pipe Command Reference.*

Table 3.1 describes the different types of connection configurations available on the BIG/ip Controller.

	NAT	SNAT	IP Forwarding	Virtual server	Forwarding virtual server
Security	Medium	High	Low (see following note)	High	High
Routable addresses required on the internal network	No	No	Yes	No	Yes
Protocols	TCP and UDP	TCP and UDP	Any IP protocol	TCP and UDP	TCP and UDP
NT Domain support	No	No	Yes	No	Yes
Active FTP support	No	Yes	Yes	Yes	Yes
Connections	Not connection-oriented	Source processing	Not connection-oriented	Destination processing only	Connection-oriented
Ports	Does not matter	Does not matter	Does not matter	Uses specific ports or wildcard	Uses specific ports or wildcard
Setup for specific nodes or hosts	Yes	Yes, but can use wildcards	No	Yes, but can use wildcard	Yes
Load balancing	No	No	No	Yes	No

Table 3.1 *Connection configuration options for the BIG/ip Controller*

◆ Note

Although IP forwarding does not require setup for specific hosts, the BIG/ip Controller supports IP filters which you can configure to restrict traffic.

Configuring virtual servers

The first step in a basic configuration is to configure virtual servers. Before you configure virtual servers, you need to know:

- If standard virtual servers or wildcard virtual servers meet the needs of your network
- Whether you need to activate optional virtual server properties

Once you know which virtual server options are useful in your network, you can:

- Define standard virtual servers
- Define wildcard virtual servers

Using standard or wildcard virtual servers

Virtual servers map to a group of content servers, firewalls, routers, or cache servers, and they are associated with one or more external interfaces on the BIG/ip Controller.

You can configure two different types of virtual servers:

- **Standard virtual servers**
  A standard virtual server represents a site, such as a web site or an FTP site, and it provides load balancing for content servers. The virtual server IP address should be the same IP address that you register with DNS for the site that the virtual server represents.
- **Wildcard virtual servers**
  A wildcard virtual server load balances transparent network devices such as firewalls, routers, or cache servers. Wildcard virtual servers are configured with an IP address of 0.0.0.0, and sometimes with a virtual port of 0.

Before you begin configuring the virtual servers, you should have the following information ready to enter:

- The IP address and service name or port number of each virtual server you are configuring

- The IP addresses and ports for each content server, firewall, cache server, or other device that the virtual servers will load balance

Note that both the F5 Configuration utility and the BIG/pipe command line utility accept host names in place of IP addresses, and also accept standard service names in place of port numbers.

Using additional features with virtual servers

After you define a virtual server, you can set up additional features, such as network address translation (NAT) or extended content verification (ECV). If you are planning on using any of these features, you may want to read the corresponding section before you actually begin the virtual server configuration process:

- Network address translations (see *Configuring NATs and IP forwarding for nodes*, on page 3-22)
- Extended Content Verification service checking (see *Configuring Extended Content Verification service checking*, on page 3-29)
- Persistence for connections that should return to the node to which they last connected (see *Configuring persistence for e-commerce and other dynamic content sites*, on page 3-36)

◆ Tip

If you prefer to use command line utilities for configuration tasks, you may find it more efficient to configure certain property settings at the time you define the virtual server. For example, to turn SSL persistence on using the ***bigpipe vip*** *command, you add the* ***special ssl*** *keyword to the end of the command when you define the virtual server.*

Defining standard virtual servers

A standard virtual server represents a specific site, such as an Internet web site or an FTP site, and it load balances content servers. The IP address that you use for a standard virtual server should match the IP address that DNS associates with the site's domain name.

◆ Note

If you are using a 3DNS Controller in conjunction with the BIG/ip Controller, the 3DNS Controller uses the IP address associated with the registered domain name in its own configuration. For details, refer to the ***3DNS Controller Administrator Guide****.*

When you define a virtual server, you actually define the virtual server at the same time that you define the nodes which the virtual server uses for load balancing. The combination of the virtual server and the group of nodes that it load balances is referred to as the ***virtual server mapping***.

To define a standard virtual server mapping in the F5 Configuration utility

1. In the navigation pane, click Virtual Servers.
2. On the toolbar, click **Add Virtual Server**.
   The Add Virtual Server screen opens.
3. In the **Address** box, enter the virtual server's IP address or host name.
4. In the **Netmask** box, type an optional netmask. If you leave this setting blank, the BIG/ip Controller uses a default netmask based on the IP address you entered for the virtual server. Use the default netmask unless your configuration requires a different netmask.
5. In the **Broadcast** box, type the broadcast address for this virtual server. If you leave this box blank, the BIG/ip Controller generates a default broadcast address based on the IP address and netmask of this virtual server.

6. In the **Port** box, either type a port number, or select a service name from the drop-down list.
7. For **Interface**, select the external (destination processing) interface on which you want to create the virtual server. Select **default** to allow the F5 Configuration utility to select the interface based on the network address of the virtual server. If no external interface is found for that network, the virtual server is created on the first external interface. If you choose **None**, the BIG/ip Controller does not create an alias and generates no ARPs for the virtual IP address. In this case, the BIG/ip Controller accepts traffic on all interfaces.
8. In Resources, click the **Node List** button.
9. In the **Node Address** box, type the IP address or host name of the first node to which the virtual server maps. If you have already defined a node address, you can choose it from the list.
10. In the **Node Port** box, type the node port number, or select the service from the drop-down list. Note that in typical applications, the node port is the same as the virtual server port. If you have already defined a node port, you can choose it from the list.
11. Click the add button (>>) to add the node to the Current Nodes list for the virtual server.
12. To add additional nodes to the virtual server mapping, type in a Node Address, Node Port, and click the add button (>>).
13. To remove nodes from the virtual server mapping, click the node listed in the Current Nodes list and click the remove button (<<).
14. After you have added or removed nodes from the Current Nodes list, click the **Add** button to save the virtual server.

To define a standard virtual server mapping on the command line

Type the **bigpipe vip** command as shown below. Note that when you define a virtual server on the command line, you can define all nodes in the mapping at once. Also, you can use host names in place of IP addresses, and that you can use standard service names in place of port numbers.

```
bigpipe vip <virt IP>:<port> define <node IP>:<port> \
   <node IP>:<port>...  <node IP>:<port>
```

For example, the following command defines a virtual server that maps to three nodes:

```
bigpipe vip 192.200.100.25:80 define 192.168.10.01:80 \
   192.168.10.02:80 192.168.10.03:80
```

Defining wildcard virtual servers

Wildcard virtual servers are a special type of virtual server designed to manage network traffic for transparent network devices, such as transparent firewalls, routers, proxy servers, or cache servers. A wildcard virtual server manages network traffic that has a destination IP address unknown to the BIG/ip Controller. A standard virtual server typically represents a specific site, such as an Internet web site, and its IP address matches the IP address that DNS associates with the site's domain name. When the BIG/ip Controller receives a connection request for that site, the BIG/ip Controller recognizes that the client's destination IP address matches the IP address of the virtual server, and it subsequently forwards the client to one of the content servers that the virtual server load balances.

However, when you are load balancing transparent nodes, a client's destination IP address is going to seem random. The client is connecting to an IP address on the other side of the firewall, router, or proxy server. In this situation, the BIG/ip Controller cannot match the client's destination IP address to a virtual server IP address. Wildcard virtual servers resolve this problem by not translating the incoming IP address at the virtual server level on the BIG/ip Controller. For example, when the BIG/ip Controller does

not find a specific virtual server match for a client's destination IP address, it matches the client's IP address to a wildcard virtual server. The BIG/ip Controller then forwards the client's packet to one of the firewalls or routers that the wildcard virtual server load balances, which in turn forwards the client's packet to the actual destination IP address.

A note about wildcard ports

When you configure wildcard virtual servers and the nodes that they load balance, you can use a wildcard port (port **0**) in place of a real port number or service name. A wildcard port handles any and all types of network services.

A wildcard virtual server that uses port **0** is referred to as a ***default wildcard virtual server***, and it handles traffic for all services. A ***port-specific wildcard virtual server*** handles traffic only for a particular service, and you define it using a service name or a port number. If you use both a default wildcard virtual server and port-specific wildcard virtual servers, any traffic that does not match either a standard virtual server or one of the port-specific wildcard virtual servers is handled by the default wildcard virtual server.

You can use port-specific wildcard virtual servers for tracking statistics for a particular type of network traffic, or for routing outgoing traffic, such as HTTP traffic, directly to a cache server rather than a firewall or router.

We recommend that when you define transparent nodes that need to handle more than one type of service, such as a firewall or a router, you specify an actual port for the node and turn off port translation for the virtual server.

◆ Note

*When you define a virtual server with port translation turned **off**, and you want to perform a service check on that node, you must configure service check intervals and timeouts using the port specified for the node. Then you can configure a service check.* *See Service checking for wildcard servers and ports, on page 3-18, for more details.*

Defining the wildcard virtual server mappings

There are two procedures required to set up a wildcard virtual server. First, you must define the wildcard virtual server. Then you must turn port translation off for the virtual server.

To define a wildcard virtual server mapping in the F5 Configuration utility

1. In the navigation pane, click Virtual Servers.
2. On the toolbar, click **Add Virtual Server**.
   The Add Virtual Server screen opens.
3. In the **Address** box, type the wildcard IP address **0.0.0.0**.
4. In the **Netmask** box, type an optional netmask.
   If you leave this box blank, the BIG/ip Controller generates a default netmask address based on the IP address of this virtual server. Use the default netmask unless your configuration requires a different netmask.
5. In the **Broadcast** box, type the broadcast address for this virtual server.
   If you leave this box blank, the BIG/ip Controller generates a default broadcast address based on the IP address and netmask of this virtual server.
6. In the **Port** box, type a port number, or select a service name from the drop-down list. Note that port **0** defines a wildcard virtual server that handles all types of services. If you specify a port number, you create a port-specific wildcard virtual server. The wildcard virtual server only handles traffic for the port specified.
7. For **Interface**, select the external (destination processing) interface on which you want to create the virtual server.
   If you choose **None**, the BIG/ip Controller does not create an alias and generates no ARPs for the virtual IP address (see the ***BIG/ip Controller Administrator Guide***, *Optimizing large configurations*, on page 2-18, for details).
8. In Resources, click the **Node List** button.

9. In the **Node Address** box, type the IP address or host name of the first node to which the virtual server maps. If you have already defined a node address, you can choose it from the list.
10. In the **Node Port** box, type the node port number, or select the service from the drop-down list. If you have already defined a node port, you can select it from the list. We recommend you choose a valid node port, such as 80, and then disabling address translation on the virtual server. This allows you to set up service checks for the node port.
11. Click the add button (>>) to add the node to the Current Nodes list for the virtual server.
12. To add additional nodes to the virtual server mapping, type in a Node Address, Node Port, and click the add button (>>).
13. To remove nodes from the virtual server mapping, click the node listed in the Current Nodes list and click the remove button (<<).
14. After you have added or removed nodes from the Current Nodes list, click the **Add** button to save the virtual server.

To turn off port translation for a wildcard virtual server in the F5 Configuration utility

After you define the wildcard virtual server with a wildcard port, you must disable port translation for the virtual server.

1. In the navigation pane, click Virtual Servers.
   The Virtual Servers screen opens.
2. In the virtual server list, click the virtual server for which you want to turn off port translation.
   The Virtual Server Properties screen opens.
3. In the Enable Translation section, clear the **Port** box.
4. Click the **Apply** button.

To define a wildcard virtual server mapping on the command line

There are two commands required to set up a wildcard virtual server. First, you must define the wildcard virtual server. Then you must turn port translation off for the virtual server. To define the virtual server, use the **bigpipe vip** command:

```
bigpipe vip <virtual IP>:<port> define <node IP>:<port> \
   <node IP>:<port>...  <node IP>:<port>
```

After you define the virtual server, you can enable or disable port translation using the following command:

```
bigpipe vip <virtual IP>:<port> translate port enable | disable
```

For example, the following command defines a wildcard virtual server that maps to three nodes. Because the nodes are firewalls and need to handle a variety of services, the virtual server is defined using port **0** (or ***** or **any**). You can specify any valid non-zero port for the node port and then turn off port translation for that port. In this example, service checks ping port 80.

```
bigpipe vip 0.0.0.0:0 define 192.168.10.01:80 \
   192.168.10.02:80 192.168.10.03:80
```

After you define the virtual server, turn off port translation for the port in the virtual server definition. In this example, port 80 is used for service checking. If you do not turn off port translation, all incoming traffic would be translated to port 80.

```
bigpipe vip 0.0.0.0:0 translate port disable
```

Allowing access to ports and services

One of the security features of the BIG/ip Controller is that all ports on the controller are locked down and unavailable for service unless you specifically open them to network access. Before clients can use the virtual servers you have defined, you must allow access to each port that the virtual servers use.

This is the second task of the three essential tasks you must complete for a basic configuration. You must perform this task after you configure virtual servers and before you configure the timer settings.

Tip

Virtual servers using the same service actually share a port on the BIG/ip Controller. This command is global, you only need to open access to a port once; you do not need to open access to a port for each instance of a virtual server that uses it.

To allow access to services in the F5 Configuration utility

Any time you create a virtual server and define a port or service with the F5 Configuration utility, the port or service is automatically enabled.

To allow access to services on the command line

Using the **bigpipe port** command, you can allow access to one or more ports at a time.

```
bigpipe port <port>... <port> enable
```

For example, in order to enable HTTP (port 80) and Telnet (port 23) services, you can enter the following **bigpipe port** command:

```
bigpipe port 80 23 enable
```

WARNING

*In order for FTP to function properly, you must allow both ports 20 and 21 (or **ftp-data** and **ftp**).*

Configuring the timer settings

Configuring timer settings is the third task of the three essential tasks you must complete for a basic configuration. You must perform this task after you configure virtual servers and after you allow access to services and ports.

There are two essential timer settings that you need to configure:

- The node ping timer defines how often the BIG/ip Controller will ping node addresses to verify whether a node is up or down. It also defines how long the BIG/ip Controller waits for a response from a node before determining that the node is unresponsive and marking the node **down**.
- The idle connection timer defines how long an inactive connection is allowed to remain open before the BIG/ip Controller deletes the record of the connection, closing it and disconnecting the client.

The service check timer is optional, and you need to set it only if you want the BIG/ip Controller to check to see if a service, or even specific content, is available on a particular node. If you plan on using simple service checks, or ECV or EAV service checks, you need to set the service check timer.

Setting the node ping timer

The node ping timer is an essential setting on the BIG/ip Controller that determines how often the BIG/ip Controller checks node addresses to see whether they are **up** and available or **down** and unavailable. The node ping timer setting applies to all nodes configured for use by the BIG/ip Controller, and it is part of the BIG/ip Controller system properties.

To set the node ping timer in the F5 Configuration utility

1. In the navigation pane, click the BIG/ip logo.
   The BIG/ip System Properties screen opens.
2. In the Node Ping section of the table, in the **Ping** box, type the frequency (in seconds) at which you want the BIG/ip Controller to ping each node address it manages. A setting of 5 seconds is adequate for most configurations.
3. In the Node Ping section of the table, in the **Timeout** box, type the number of seconds you want the BIG/ip Controller to wait to receive a response to the ping. If the BIG/ip Controller does not receive a response to the ping before

the node ping timeout expires, the BIG/ip Controller marks the node **down** and does not use it for load balancing. A setting of 16 seconds is adequate for most configurations.

To set the node ping timer on the command line

To define node ping settings, you use two commands. First, you set the node ping frequency using the **bigpipe tping_node** command, and then you set the node ping timer using the **bigpipe timeout_node** command.

```
bigpipe tping_node <seconds>
bigpipe timeout_node <seconds>
```

For example, the following commands sets the ping frequency at 5 seconds, and the timer to 16 seconds, which should be adequate for most configurations.

```
bigpipe tping_node 5
bigpipe timeout_node 16
```

Setting the timer for reaping idle connections

The BIG/ip Controller supports two timers for reaping idle connections, one for TCP traffic and one for UDP traffic. These timers are essential, and if they are set too high, or not at all, the BIG/ip Controller may run out of memory. Each individual port on the BIG/ip Controller has its own idle connection timer settings.

WARNING

The BIG/ip Controller accepts UDP connections only if you set the UDP idle connection timer.

To set the inactive connection timer in the F5 Configuration utility

1. In the navigation pane, click the plus (+) next to Virtual Servers.
   The Virtual Server tree opens and displays the Ports option.

2. Click Ports.
   The Global Virtual Ports screen opens.
3. In the **Port** box, click the port number or service name for which you want to configure the idle connections timeouts.
   The Global Virtual Port screen opens.
4. In the **Idle Connection Timeout TCP** box, type the number of seconds you want to elapse before the BIG/ip Controller drops an idle TCP connection. For HTTP connections, 60 seconds should be adequate, but for other services such as Telnet, higher settings may be necessary.
5. In the **Idle Connection Timeout UDP** box, type the number of seconds you want to elapse before the BIG/ip Controller drops UDP connections.
6. Click the **Apply** button.

To set TCP idle connection timers on the command line

Use the **bigpipe treaper** to define a TCP idle connection timeout for one or more ports at a time. For HTTP connections we recommend only 60 seconds, but for other services such as Telnet we recommend higher settings. The default setting for this timer is 16 minutes (1005 seconds). Use the following syntax for this command.

```
bigpipe treaper <port>... <port> <seconds>
```

For example, the following command sets a 120 second time limit for idle connections on port 443:

```
bigpipe treaper 443 120
```

To set UDP idle connection timers on the command line

You can define a UDP idle connection timeout for one or more ports at a time using the **bigpipe udp** command.

```
bigpipe udp <port>... <port> <seconds>
```

For example, the following command sets a 120-second time limit for idle connections on port 53:

```
bigpipe udp 53 120
```

Setting the service check timer

The service check feature is similar to node ping, but instead of testing the availability of a server, it tests the availability of a particular service running on a server. The service check timer affects the three different types of service checks: simple service check, ECV service check, and EAV service check. To set up simple service check, you need only set the service check timer as described below. To set up ECV service check or EAV service check, however, you need to configure additional settings (see *Configuring Extended Content Verification service checking*, on page 3-29).

Note that each individual service managed by the BIG/ip Controller has its own service check timer settings.

To set the service check timer in the F5 Configuration utility

1. In the navigation pane, click the plus (+) next to Nodes.
   The Nodes tree opens and displays the Ports option.
2. Click Ports.
   The Global Node Ports screen opens.
3. Click the port you want to configure.
   The Global Node Port Properties screen opens.
4. In the **Frequency** box, type the frequency (in seconds) at which you want the BIG/ip Controller to check the service on the node for all defined nodes using this port. Five seconds is adequate for most configurations.
5. In the **Timeout** box, type the number of seconds you want the BIG/ip Controller to wait to receive a response to the service check. If the BIG/ip Controller does not receive a response to the service check before the timeout expires, the BIG/ip Controller marks the service on the node **down** and does not use it for load balancing. Sixteen (16) seconds is adequate for most configurations.
6. Click the **Apply** button.

To set the service check timer on the command line

To define service check settings, you actually use a series of two commands. First, you set the service check frequency using the **bigpipe tping_svc** command, and then set the service check timer using the **bigpipe timeout_svc** command.

```
bigpipe tping_svc <port> <seconds>
bigpipe timeout_svc <port> <seconds>
```

For example, the following series of commands sets the service check frequency at 5 seconds, and the timer to 16 seconds, which is adequate for most configurations.

```
bigpipe tping_svc 80 5
bigpipe timeout_svc 80 16
```

Service checking for wildcard servers and ports

When you configure a wildcard virtual server with a **0** port using nodes with standard ports, such as 80, with port translation turned **off**, the BIG/ip Controller uses the standard service check timeout values (e.g., port 80) to service check the port. For more information about setting the service check timer, see *Setting the service check timer*, on page 3-17

Using the simple keyword

The simple keyword is being phased out in future releases. This information is provided in order to support existing configurations.

The **simple** keyword is necessary only if you specified a node port of **0**. In previous versions of the BIG/ip Controller, this was the only way to set up a wildcard virtual server that handled connections for all services. However, we now recommend that you specify a node port and then turn off port translation for the virtual server.

To set up a simple service check for this type of virtual server, add the following entry to the **/etc/bigd.conf** file. Use the following syntax to set a check on a node where the check port is not the node port:

simple [<node addr>:]<node port> <check port>

For example, a wildcard server is defined with a wildcard port, like this:

bigpipe vip 0.0.0.0:0 define n1:0

In this case, you must use the **simple** keyword to designate the wildcard <**server:**><**port**> and <**check port**> for the service check:

simple n1:0 80

Changing the global load balancing mode

Changing the global load balancing mode is one of the optional tasks you can perform after you have completed the three main tasks of a basic configuration. This means you already have:

- Configured virtual servers
- Configured access to ports and services
- Configured the timer settings

After you complete the basic tasks, you can change the global load balancing mode. The default global load balancing mode on the BIG/ip Controller is Round Robin, and it simply passes each new connection request to the next server in line, eventually distributing connections evenly across the array of machines being load balanced. Round Robin mode works well in most configurations, especially if the equipment that you are load balancing is roughly equal in processing speed and memory. If you want to use the Round Robin load balancing mode, you can skip this section, and begin configuring features that you want to add to the basic configuration.

However, if you are working with servers that differ significantly in processing speed and memory, you may want to switch to Ratio load balancing mode. In Ratio mode, the BIG/ip Controller distributes connections among machines according to ratio weights

that you define, where the number of connections that each machine receives over time is proportionate to the ratio weight you define for each machine.

Tip

*The default ratio weight for a node is **1**. If you keep the default ratio weight for each node in a virtual server mapping, the nodes receive an equal proportion of connections as though you were using Round Robin load balancing.*

Note

*The BIG/ip Controller also supports more advanced dynamic load balancing modes that may be suitable for your site. Refer to the **BIG/ip Controller Administrator Guide**, Working with Special Features, for more information about working with specialized load balancing modes.*

Using Ratio mode

If you want to switch from Round Robin to Ratio mode, you need to complete two separate configuration tasks:

- Define the ratio weight (percentage of connections to be handled) for each individual node address.
- Set the load balancing mode to Ratio as described in the following section.

Setting ratio weights for node addresses in a virtual server mapping

Once you switch to Ratio load balancing mode, you need to set the ratio weight for each node address. Weights are a property of a node's IP address, and the default ratio weight for any given node address is **1**.

To set ratio weights in the F5 Configuration utility

1. In the navigation pane, click Nodes.

2. In the Nodes list, select the node for which you want to set the ratio weight.
   The Node Properties screen opens.
3. In the Node Properties screen, select the **Address** of the node.
   The Global Node Address Properties screen opens.
4. In the **Ratio or Priority** box, replace the default ratio weight with the ratio weight of your choice.
5. Click the **Apply** button to save your changes.

To set ratio weights on the command line

The **bigpipe ratio** command sets the ratio weight for one or more node addresses:

```
bigpipe ratio <node IP>... <node IP> <ratio weight>
```

The following example defines ratio weights for three node addresses. The first command sets the first node to receive half of the connection load. The second command sets the two remaining node addresses to each receive one quarter of the connection load.

```
bigpipe ratio 192.168.10.01  2
bigpipe ratio 192.168.10.02  192.168.10.03  1
```

Switching to Ratio mode

First, you should set the load balancing mode to Ratio. The load balancing mode is actually a property of the BIG/ip Controller system, and it applies to all virtual servers defined on the system.

To switch the system to Ratio mode in the F5 Configuration utility

1. In the navigation pane, click Virtual Servers.
   The Virtual Servers screen opens.
2. In the **Node List Load Balancing Method** box, select **Ratio (node address)**.

To switch the system to Ratio mode on the command line

To switch the system to Ratio mode, type the following BIG/pipe command:

```
bigpipe lb ratio
```

Configuring NATs and IP forwarding for nodes

Configuring NATs and IP forwarding are optional tasks you can configure after you have completed the three main tasks of a basic configuration. This means you already have:

- Configured virtual servers
- Configured access to ports and services
- Configured the timer settings

After you complete the basic tasks, you can configure network address translation and IP forwarding on the BIG/ip Controller.

The IP addresses that identify nodes on the BIG/ip Controller's internal network need not be routable on the external network. This protects nodes from illegal connection attempts, but it also prevents nodes (and other hosts on the internal network) from receiving direct administrative connections, or from initiating connections to clients, such as mail servers or databases, on the BIG/ip Controller's external interface (destination processing).

Using network address translation resolves this problem. Network address translations (NATs) assign to a particular node a routable IP address that the node can use as its source IP address when connecting to servers on the BIG/ip Controller's external interface. You can use the NAT IP address to connect directly to the node through the BIG/ip Controller, rather than having the BIG/ip Controller send you to a random node according to the load

balancing mode. IP forwarding provides functionality similar to a NAT. If your network does not support NATs, you may want to consider using IP forwarding.

◆ Note

In addition to these options, you can set up forwarding virtual servers which allow you to selectively forward traffic to specific addresses. The BIG/ip Controller maintains statistics for forwarding virtual servers. For more information about forwarding virtual servers, see the ***BIG/ip Controller Administrator Guide****, Working with Special Features.*

There are three configuration options on the BIG/ip Controller that you can use to control network access, and you need to identify which method is suitable for your needs:

- **Network Address Translation (NAT)**
  A network translation address provides a routable alias IP address that a node can use as its source IP address when making or receiving connections to clients on the external network. You can configure a unique NAT for each node address included in a virtual server mapping.
  Note that NATs do not support port translation, and are not appropriate for FTP. You cannot define a NAT if you configure a default SNAT.
- **Secure Network Address Translation (SNAT)**
  A secure network address translation provides functionality similar to that of firewalls. A SNAT defines a routable alias IP address that one or more nodes can use as a source IP address only when making connections to hosts on the external network. SNAT addresses support port translation, and they also prevent hosts on the external network from connecting directly to the node.
  Note that SNAT only supports TCP and UDP. SNAT also features support for both passive and active FTP. You cannot define a NAT if you define a default SNAT.
- **IP forwarding**
  IP forwarding does not translate node addresses. Instead, it simply exposes the node's IP address to the BIG/ip Controller's external network and clients can use it as a standard routable

address. When you turn IP forwarding on, the BIG/ip Controller acts as a router when it receives connection requests for node addresses. IP forwarding itself does not provide security features, but you can use the IP filter feature to implement a layer of security (see *Setting up IP forwarding*, on page 3-28) that can help protect your nodes.

WARNING

NATs and SNATs do not support the NT Domain or CORBA protocols. Instead of using NATs or SNATs, you need to configure IP forwarding (see Setting up IP forwarding, on page 3-28).

Defining a standard network address translation (NAT)

When you define standard network address translations (NATs), you need to create a separate NAT for each node that requires a NAT. You also need to use unique IP addresses for NAT addresses; a NAT IP address cannot match an IP address used by any virtual or physical servers in your network. You can configure a NAT with the F5 Configuration utility or from the command line.

To configure a NAT in the F5 Configuration utility

1. In the navigation pane, click NATs.
   The Network Address Translations screen opens.
2. On the toolbar, click **Add NAT**.
   The Add Nat screen opens.
3. In the **Node Address** box, type the IP address of the node.
4. In the **NAT Address** box, type the IP address that you want to use as the node's alias IP address.
5. In the **NAT Netmask** box, type an optional netmask. If you leave this box blank, the BIG/ip Controller generates a default broadcast address based on the IP address and netmask of this virtual server.

6. In the **NAT Broadcast** box, type the broadcast address. If you leave this box blank, the BIG/ip Controller generates a default broadcast address based on the IP address and netmask of this NAT.
7. In the **Interface** box, you can select an external interface (destination processing) on which the NAT address is to be used. Note that this setting only applies if the BIG/ip Controller has more than one external interface.
8. Click the **Apply** button.

To configure a NAT on the command line

The **bigpipe nat** command defines one NAT for one node address.

```
bigppipe nat <node addr> to <NAT addr>
```

Defining a secure network address translation (SNAT)

When you define secure network address translations (SNATs), you can assign a single SNAT address to multiple nodes. Note that a SNAT address does not necessarily have to be unique; for example, it can match the IP address of a virtual server.

SNAT addresses have global properties that apply to all SNATs that you define in the BIG/ip Controller configuration as well as to the SNAT mappings you define. You can configure SNATs in the F5 Configuration utility or from the command line.

Setting SNAT global properties

The SNAT feature supports three global properties that apply to all SNAT addresses:

- **Connection limits**
  The connection limit applies to each node that uses a SNAT, and each individual SNAT can have a maximum of 50,000 simultaneous connections.

- **TCP idle connection timeout**
  This timer defines the number of seconds that TCP connections initiated using a SNAT address are allowed to remain idle before being automatically disconnected.
- **UDP idle connection timeout**
  This timer defines the number of seconds that UDP connections initiated using a SNAT address are allowed to remain idle before being automatically disconnected. This value should not be set to **0**.

To configure SNAT global properties in the F5 Configuration utility

1. In the navigation pane, click Secure NATs.
   The Secure Network Address Translations screen opens.
2. In the **Connection Limit** box, type the maximum number of connections you want to allow to each node using a SNAT. To turn connection limits off, set the limit to **0**. If you turn connection limits on, keep in mind that each SNAT can support only 50,000 simultaneous connections.
3. In the **TCP Idle Connections** box, type the number of seconds that TCP connections initiated by a node using a SNAT are allowed to remain idle.
4. In the **UDP Idle Connections** box, type the number of seconds that UDP connections initiated by a node using a SNAT are allowed to remain idle. This value should not be set to **0**.
5. Click the **Apply** button.

To configure SNAT global properties on the command line

Configuring global properties for a SNAT requires that you enter three **bigpipe** commands. The following command sets the maximum number of connections you want to allow for each node using a SNAT.

```
bigpipe snat limit <value>
```

The following commands set the TCP and UDP idle connection timeouts:

```
bigpipe snat timeout tcp <seconds>
bigpipe snat timeout udp <seconds>
```

Configuring SNAT address mappings

Once you have configured the SNAT global properties, you can configure SNAT address mappings. The SNAT address mappings define each SNAT address, and also define the node or group of nodes that uses the SNAT address. Note that a SNAT address does not necessarily have to be unique; for example, it can match the IP address of a virtual server. A SNAT address cannot match an address already in use by a NAT or another SNAT address.

To configure a SNAT mapping in the F5 Configuration utility

1. In the navigation pane, click Secure NATs.
   The Secure Network Address Translations screen opens.
2. On the toolbar, click **Add SNAT**.
   The Add SNAT screen opens.
3. In the **Translation Address** box, type the IP address that you want to use as the alias IP address for the node(s).
4. In the **Interface** box, you can select the external interface (destination processing) on which the SNAT address is to be used. Note that this setting applies only if your BIG/ip Controller has more than one destination processing interface.
5. In the **Original Address** box, type the IP address of the node or nodes that are assigned to the SNAT. Click the add button (>>) to add the address to the **Current List**.
6. To remove an address from the **Current List**, click the remove button (<<).
7. Click the **Apply** button.

To configure a SNAT mapping on the command line

The **bigpipe snat** command defines one SNAT for one or more node addresses.

```
bigpipe snat map <node addr>... <node addr> to <SNAT addr>
```

For example, the command below defines a secure network address translation for two nodes:

```
bigpipe snat map 192.168.75.50 192.168.75.51 to 192.168.100.10
```

Setting up IP forwarding

If you do not want to translate addresses with a NAT or SNAT, you can use the IP forwarding configuration option. IP forwarding is an alternate way of allowing nodes to initiate or receive direct connections from the BIG/ip Controller's external network. IP forwarding exposes all of the node IP addresses to the external network, making them routable on that network. If your network uses the NT Domain or CORBA protocols, IP forwarding is an option for direct access to nodes.

To set up IP forwarding, you need to complete two tasks:

- **Turn IP forwarding on**
  The BIG/ip Controller uses a system control variable to control IP forwarding, and its default setting is **off**.
- **Verify the routing configuration**
  You probably have to change the routing table for the router on the BIG/ip Controller's external network. The router needs to direct packets for nodes to the BIG/ip Controller, which in turn directs the packets to the nodes themselves.

Turning on IP forwarding

IP forwarding is a property of the BIG/ip Controller system, and it is controlled by the system control variable **net.inet.ip.forwarding**.

To set the IP forwarding system control variable in the F5 Configuration utility

1. In the navigation pane, click the BIG/ip logo.
   The BIG/ip System Properties screen opens.
2. On the toolbar, click **Advanced Properties**.
   The BIG/ip System Control Variables screen opens.
3. Check the **Allow IP Forwarding** box.
4. Click the **Apply** button.

To set the IP forwarding system control variable on the command line

Use the standard **sysctl** command to set the variable. The default setting for the variable is **0**, which is **off**. You want to change the setting to **1**, which is **on**:

```
sysctl -w net.inet.ip.forwarding=1
```

To permanently set this value, you can use a text editor, such as **vi** or **pico**, to manually edit the **/etc/rc.sysctl** file. For additional information about editing this file, see the ***BIG/ip Controller Reference Guide***, *BIG/ip Controller System Control Variables*.

Addressing routing issues for IP forwarding

Once you turn on IP forwarding, you probably need to change the routing table on the default router. Packets for the node addresses need to be routed through the BIG/ip Controller. For details about changing the routing table, refer to your router's documentation.

Configuring Extended Content Verification service checking

Extended content verification service checking is another feature you can configure after you have performed the three basic configuration tasks. Extended content verification service check is a special type of service check that actually retrieves content from a

server. If the content matches the expected result, the BIG/ip Controller marks the node **up** and uses it for load balancing. If the content does not match, or if the server does not return content, the BIG/ip Controller marks the node **down**, and does not use it for load balancing.

You can set up ECV service check in the F5 Configuration utility, or you can use a text editor, such as **vi** or **pico**, to manually create the **/etc/bigd.conf** file, which stores ECV information.

ECV service check is most frequently used to verify content on web servers, although you can use it for more advanced applications, such as verifying firewalls or mail servers. This section focuses on setting up ECV for web servers. For details about using advanced ECV service check options, see the ***BIG/ip Controller Administrator Guide***, *Working with Advanced Service Check Options*.

◆ Note

It is important to note that the intervals and timeouts for service checks apply to EAV and ECV service checks. These timeouts are configured by setting the service check timers. For more information about setting these timers, see Configuring the timer settings, on page 3-13.

ECV service check properties

ECV service check is a property of both a node port and a node. If you define ECV service check settings for a node port, all nodes that use the port inherit the ECV service check settings. You can override these settings by defining ECV service check settings for the node itself.

There are actually three different types of ECV service check settings that you can define:

- **ECV normal**
  An ***ECV normal*** service check requires that the BIG/ip Controller mark a node **up** (available for load balancing) when the retrieved content matches the expected result. For example, if the home page for your web site included the words **Welcome**

home, you could set up an ECV service check to look for the string **"Welcome home"**. A match for this string would mean that the web server is up and available.

- **ECV SSL**
  An ***ECV SSL*** service check performs the same function as an ECV normal service check, but it is designed to work with secure servers that use the SSL protocol, rather than standard servers using HTTP. The BIG/ip Controller uses SSL version 3, as do popular web browsers, but it is backward-compatible for web servers that support only version 2.
- **ECV reverse**
  In contrast, an ***ECV reverse*** service check requires that the BIG/ip Controller mark a node **down** (not available for load balancing) when the retrieved content matches the expected result. For example, if the content on your web site home page is dynamic and changes frequently, you may prefer to set up a reverse ECV service check that looks for the string **"Error"**. A match for this string would mean that the web server was down.

WARNING

When the BIG/ip Controller checks content looking for a match, it reads through the content until the service check times out, or until the read reaches 5,000 bytes, whichever comes first. When you choose text, an HTML tag, or an image name to search for, be sure to pick one that appears in the first 5,000 bytes of the web page.

Writing regular expressions for ECV service checks

When you set up an ECV service check for a web server, you need to define a send string and a receive expression. A ***send string*** is the request that the BIG/ip Controller sends to the web server. Send strings typically request that the server return a specific web page, such as the default page for a web site. For example, the most common send string is **"GET /"** which simply retrieves the default HTML page for a web site. The ***receive expression*** is the text string that the BIG/ip Controller looks for in the returned web page.

Receive expressions use regular expression syntax, and they are not case-sensitive. Although regular expressions can be complex, you will find that simple regular expressions are adequate for most ECV service checks.

The corresponding receive string could be any simple text string included in your home page, such as text, HTML tags, or image names.

Sample send strings

The send string below is probably the most common send string, and it retrieves the default HTML page for a web site. Note that all send strings are enclosed by quotation marks (" ") inside the **/etc/bigd.conf** file.

```
"GET /"
```

To retrieve a specific page from a web site, simply enter a fully qualified path name:

```
"GET /www/support/customer_info_form.html"
```

Sample receive expressions

The most common receive expressions contain a text string that would be included in a particular HTML page on your site. The text string can be regular text, HTML tags, or image names. Note that all receive expressions are enclosed by quotation marks (" ").

For example, the following receive expression attempts to match the text **Welcome**, and it is useful for ECV reverse service checks:

```
"welcome"
```

The sample receive expression below searches for a standard HTML tag. Note that even though you are searching for an HTML tag, you still need to enclose the regular expression with quotation marks (" ").

```
"<HEAD>"
```

You can also use null receive expressions, formatted as the one shown below. When you use a null receive expression, the BIG/ip Controller considers any content retrieved to be a match.

```
""
```

Null receive expressions are suitable only for ECV normal and ECV SSL. Note, however, that if you use them you run the risk of the BIG/ip Controller considering an HTML error page to be a successful service check.

◆ Note

*The regular expression syntax discussed here is not the same as the **wildcard syntax** that is commonly used in command shells. For more information about regular expression, see the man page for **re_format**. To view the man page for **re_format**, type **man re_format** at the command line.*

Setting up ECV service check in the F5 Configuration utility

In the F5 Configuration utility, you can set ECV service check options in the Global Node Port Properties screen, and also in individual Node Properties screens. Regardless of which screen you use to configure the options, the steps are the same.

To set up ECV service check in the F5 Configuration utility

1. In the navigation pane, click Nodes.
   The Nodes screen opens.
2. Select a node from the list.
   The Node Properties screen opens.
3. If you want to configure ECV service check options, stay in this screen. If you want to configure ECV service check options for the port that the node uses, click the port number listed next to the IP address of the node.
4. Click the **ECV** button.
5. In the **Type** box, choose the type of ECV service check you want to set up: normal, reverse, or SSL.

6. In the **Send String** box, type the send string that requests the web page. Note that the F5 Configuration utility automatically places quotation marks around the string itself. For example, the following string retrieves the default HTML page for the site.

   ```
   GET /
   ```

7. In the **Receive Rule** box, type the receive expression that the BIG/ip Controller should look for in the returned web page. For example, the following receive expression looks for a text string in a web page:

   ```
   Welcome home!
   ```

8. Click the **Apply** button.

Manually configuring and testing the /etc/bigd.conf file

You can set up ECV service check on the command line by creating an **/etc/bigd.conf** file in a text editor such as **vi** or **pico**. Each line in the **/etc/bigd.conf** file defines a send string and a receive expression for one node, or for one port. Remember that when you define a ECV service check for a port, all nodes that use the port inherit the service check settings.

Changes to the **/etc/bigd.conf** do not take effect until the system is rebooted, or **bigd** is restarted. To restart **bigd**, simply run the command **bigd**.

The BIG/ip Controller provides a command line tool that allows you to verify the syntax of the **/etc/bigd.conf** file before the system begins using it. Once you set up the file, we recommend that you test it before you reboot the system or restart the **bigd** daemon and begin using the file.

Setting up the /etc/bigd.conf file

The **/etc/bigd.conf** file uses three different types of syntax for lines in the file that correspond to the three different types of service check that you can configure: ECV normal, ECV SSL, and ECV reverse. The following sections describe the syntax for each type, and provide some useful examples.

To set up an ECV normal service check

The line for a normal ECV service check begins with the keyword **active**. The <**node IP**> parameter is optional, and you need to include it only if you are defining ECV service check for a specific node.

```
active [<node IP>:]<port> "<send_string>" "<recv_expr>"
```

For example, the following line sets up a normal ECV service check for a node, where the BIG/ip Controller looks for the text **Welcome** in the default page for the site.

```
active 192.168.100.10:80 "GET /" "welcome"
```

To set up ECV SSL service check

The line for an SSL ECV service check begins with the keyword **ssl**. The <**node IP**> parameter is optional, and you need to include it only if you are defining ECV service check for a specific node.

```
ssl [<node IP>:]<port> "<send_string>" "<recv_expr>"
```

For example, the following line sets up an SSL ECV service check for a node port. Note that the receive expression is null. When you use a null receive expression, the BIG/ip Controller considers any retrieved content to be a match.

```
ssl 443 "GET /www/orders/order_form.html" ""
```

To set up ECV reverse service check

The line for a reverse ECV service check begins with the keyword **reverse**. The <**node IP**> parameter is optional, and you need to include it only if you are defining ECV service check for a specific node.

```
reverse [<node IP>:]<port> "<send_string>" "<recv_expr>"
```

For example, the following line sets up a reverse ECV service check for a node port. Note that the receive expression is null. When you use a null receive expression, the BIG/ip Controller considers any retrieved content to be a match.

```
reverse 80 "GET /" ""
```

Testing /etc/bigd.conf syntax

To test /etc/bigd.conf syntax

You can test your ECV syntax in the **bigd.conf** file using the following **bigd** command:

```
/sbin/bigd -d
```

This command parses the file, checks ECV syntax, reports any errors, and then exits.

◆ Note

*The **/etc/bigd.conf** file is read once at startup. If you change the file on the command line, you must reboot or restart **bigd** for the changes to take effect. If you make changes in the F5 Configuration utility, clicking the **Apply** button makes changes and restarts **bigd**. For more information about **bigd**, see the **BIG/ip Controller Reference Guide**, System Utilities.*

Configuring persistence for e-commerce and other dynamic content sites

Persistence is another feature you can configure after you have completed the three main tasks of a basic configuration. This means you already have:

- Configured virtual servers
- Configured access to ports and services
- Configured the timer settings.

If you are setting up an e-commerce or other type of dynamic content site, you may need to configure persistence on the BIG/ip Controller. Whether you need to configure persistence or not simply depends on how you store client-specific information, such as items in a shopping cart, or airline ticket reservations. For example, you may store the airline ticket reservation information in

a back-end database that all nodes can access; or on the specific node to which the client originally connected; or in a cookie on the client's machine.

If you store client-specific information on specific nodes, you need to configure persistence. When you turn on persistence, returning clients can bypass load balancing and instead can go to the node where they last connected in order to get to their saved information.

The BIG/ip Controller tracks information about individual persistent connections, and keeps the information only for a given period of time. The way in which persistent connections are identified depends on the type of persistence. The BIG/ip Controller supports two basic types of persistence:

- **SSL persistence**
  SSL persistence is a type of persistence that tracks SSL connections using the SSL session ID, and it is a property of individual virtual servers. Using SSL persistence can be particularly important if your clients typically have translated IP addresses or dynamic IP addresses, such as those that Internet service providers typically assign. Even when the client's IP address changes, the BIG/ip Controller still recognizes the connection as being persistent based on the session ID.
- **Simple persistence**
  Simple persistence supports TCP and UDP protocols, and it tracks connections based only on the client IP address. When a client requests a connection to a virtual server that supports simple persistence, the BIG/ip Controller checks to see if that client previously connected, and if so, returns the client to the same node.

You may want to use SSL persistence and simple persistence together. In situations where an SSL session ID times out, or where a returning client does not provide a session ID, you may want the BIG/ip Controller to direct the client to the original node based on

the client's IP address. As long as the client's simple persistence record has not timed out, the BIG/ip Controller can successfully return the client to the appropriate node.

Note

The BIG/ip Controller also supports several advanced persistence modes. For more information about these advanced modes, see Working with Advanced Persistence Options, in the ***BIG/ip Controller Administrator Guide****.*

Setting up SSL persistence

SSL persistence is a property of a virtual server, and to set it up for a given virtual server, you need to do two things:

- Turn SSL persistence on.
- Set the SSL session ID timeout, which determines how long the BIG/ip Controller stores a given SSL session ID before removing it from the system.

To configure SSL persistence in the F5 Configuration utility

1. In the navigation pane, click Virtual Servers.
   The Virtual Server screen opens and displays each virtual server's IP address and associated port number or service name in the Current List.
2. Click the appropriate virtual server in the list.
   The Virtual Server Properties screen opens.
3. In the **Application Persistence** box, click **None**. If you have already configured an application persistence method, the name of the method displays in this box.
   The Virtual Server Application Persistence screen opens.
4. Click the **SSL Persistence** button.
5. In the **Timeout** box, type the number of seconds that the BIG/ip Controller should store SSL sessions IDs before removing them from the system.

6. Click the **Apply** button.

To configure SSL persistence on the command line

On the command line, you actually set SSL persistence properties at the time you define the virtual server.

```
bigpipe vip <virt addr>:<port> define <node addr>:<port> \
   special ssl <persist time>
```

For example, the following command sets SSL persistence for a virtual server, where the session ID record is valid for one hour (3600 seconds).

```
bigpipe vip v1:ssl define n1:ssl n2:ssl special ssl 3600
```

Setting up simple persistence

You can set simple persistence properties for both an individual virtual server, and for a port. Individual virtual server persistence settings can override those of the port. When you set simple persistence on a port, all virtual servers that use the given port inherit the port's persistence settings.

Setting simple persistence on virtual servers

Persistence settings for virtual servers apply to both TCP and UDP persistence. Note that the persistence settings for the virtual server override the persistence settings defined for the port that the virtual server uses. When the persistence timer is set to a value greater than 0, persistence is **on**. When the persistence timer is set to 0, persistence is **off**.

To configure simple persistence for virtual servers in the F5 Configuration utility

1. In the navigation pane, click the plus Virtual Servers.
   The Virtual Server screen opens.
2. Select the virtual server for which you want to configure simple persistence.
   The Virtual Server Properties screen opens.

3. In the Persistence section, type the following information:
   - **Timeout (seconds)**
     Set the number of seconds for persistence on the virtual server. This value overrides the persistence timeout set at the port level. (This option is not available if you are using rules.)
   - **Mask**
     Set the persistence mask for the virtual server. The persistence mask determines persistence based on the portion of the client's IP address that is specified in the mask.
4. Click the **Apply** button.

To configure simple persistence for virtual servers on the command line

The **bigpipe vip** command sets simple persistence for one or more virtual servers. Note that a timeout greater than 0 turns persistence **on**, and a timeout of 0 turns persistence **off**.

```
bigpipe vip <virt IP>:<port> persist <timeout>
```

After you set the persistence timeout, you can specify a netmask. To set a persist mask, use the following syntax:

```
bigpipe vip <virt IP>:<port> persist mask <mask IP>
```

For example, the first command sets simple persistence for the virtual server, where the persistent connection information expires after one hour (3600 seconds). The second command sets a persist mask of 255.255.0.0.

```
bigpipe vip 192.168.100.10:80 persist 3600
bigpipe vip 192.168.100.10:80 persist mask 255.255.0.0
```

Setting simple persistence on ports

Defining persistence on a port requires only that you set the persistence timer; you do not actually have to turn the persistence **on** and **off**. When the persistence timer is set to 0, persistence is **off**. When it is set to a value greater than 0, persistence is **on**.

To configure simple persistence for ports in the F5 Configuration utility

1. In the navigation pane, click the plus (+) next to Virtual Servers.
   The Virtual Server tree opens and displays the Ports option.
2. Click Ports.
   The Global Virtual Ports screen opens.
3. In the **Port** box, select the port number or service name for which you want to configure the simple persistence.
   The Global Virtual Port screen opens.
4. In the **Simple Persistence (seconds)** box, type the number of seconds that you want the BIG/ip Controller to store persistent connection information.
5. Click the **Apply** button.

To configure simple persistence for ports on the command line

The **bigpipe persist** command sets simple persistence for one or more ports.

```
bigpipe persist <port>... <port> <seconds>
```

For example, the following command sets simple persistence for port 80, where the persistent connection information expires after one hour (3600 seconds).

```
bigpipe persist 80 3600
```

Configuring and synchronizing redundant systems

Another feature you can configure after you have performed the three basic configuration tasks is configuration synchronization.

Redundant BIG/ip Controller systems have special settings that you need to configure, such as interface fail-safe settings. One convenient aspect of configuring a redundant system is that once

you have configured one of the controllers, you can simply copy the configuration to the other controller in the system using the configuration synchronization feature in the **bigpipe** command line tool or in the F5 Configuration utility.

There are two basic aspects about working with redundant systems:

- Synchronizing configurations between two controllers
- Configuring fail-safe settings for the interfaces

Preparing to use the synchronization command

Before you can use the **bigpipe configsync** command or the F5 Configuration utility to synchronize domestic HA redundant BIG/ip Controllers, you must first run the **config_failover** command. This command performs the following tasks:

- Checks for a fail-over IP address for the other controller in BIG/db.
- Verifies that the **AllowHosts** entry in the **/etc/sshd_config** file includes the IP address of the other controller in the redundant configuration.
- Runs the **ssh-keygen** command which creates the security keys for the controller.
- Shares the security keys with the other controller in the redundant system.

To run the **config_failover** command, type the following command from the command line:

```
config_failover
```

The **config_failover** utility prompts you for the root password of the other controller in the redundant system before it generates the security keys for the BIG/ip Controller.

Synchronizing configurations between controllers

Once you complete the initial configuration on the first controller in the system, you can synchronize the configurations between the active unit and the standby unit. When you synchronize a configuration, the following configuration files are copied to the other BIG/ip Controller:

- **The common keys in BIG/db**
- **/etc/bigip.conf**
  The **/etc/bigip.conf** file stores virtual server and node definitions and settings, including node ping settings, the load balancing mode, and NAT and SNAT settings.
- **/etc/bigd.conf**
  The **/etc/bigd.conf** file stores service check settings.
- **/etc/hosts.allow**
  The **/etc/hosts.allow** file stores the IP addresses that are allowed to make administrative shell connections to the BIG/ip Controller.
- **/etc/hosts.deny**
  The **/etc/hosts.deny** file stores the IP addresses that are not allowed to make administrative shell connections to the BIG/ip Controller.
- **User account files**
- **/etc/ipfw.conf** and **/etc/ipfw.filt**
  The **/etc/ipfw.conf** and **/etc/ipfw.filt** files store IP filter settings.
- **rc.sysctl**
  The **rc. sysctl** file contains system control variable settings.
- **/etc/rateclass.conf**
  The **/etc/rateclass.conf** file stores rate class definitions.
- **/etc/ipfwrate.conf** and **/etc/ipfwrate.filt**
  The **/etc/ipfwrate.conf** and **/etc/ipfwrate.filt** files store IP filter settings for filters that also use rate classes.
- **/etc/snmpd.conf**
  The **/etc/snmpd.conf** file stores SNMP configuration settings.

If you use command line utilities to set configuration options, be sure to save the current configuration to the file before you use the configuration synchronization feature. Use the following **bigpipe** command to save the current configuration:

```
bigpipe -s
```

WARNING

If you are synchronizing with a controller that already has configuration information defined, we recommend that you back up that controller's original configuration file(s).

To synchronize the configuration in the F5 Configuration utility

1. In the navigation pane, click the BIG/ip logo.
   The BIG/ip System Properties screen opens.
2. On the toolbar, click the **Sync Configuration** button.
   The Sync Configuration screen opens.
3. Click the **Synchronize** button.

To synchronize the configuration on the command line

You use the **bigpipe configsync** command to synchronize configurations. When you include the **all** option in the command, all the configuration files are synchronized between machines.

```
bigpipe configsync all
```

If you want to synchronize only the **/etc/bigip.conf** file, you can use the same command without any options:

```
bigpipe configsync
```

Configuring fail-safe settings

For maximum reliability, the BIG/ip Controller supports failure detection on both internal and external interface cards. When you arm the fail-safe option on an interface card, the BIG/ip Controller monitors network traffic going through the interface. If the BIG/ip Controller detects a loss of traffic on an interface when half of the

fail-safe timeout has elapsed, it attempts to generate traffic. An interface attempts to generate network traffic by issuing ARP requests to nodes accessible through the interface. Also, an ARP request is generated for the default route if the default router is accessible from the interface. Any traffic through the interface, including a response to the ARP requests, averts a fail-over.

If the BIG/ip Controller does not receive traffic on the interface before the timer expires, it initiates a fail-over, switches control to the standby unit, and reboots.

WARNING

You should arm the fail-safe option on an interface only after the BIG/ip Controller is in a stable production environment. Otherwise, routine network changes may cause fail-over unnecessarily.

Arming fail-safe on an interface

Each interface card installed on the BIG/ip Controller has a unique name, which you need to know when you set the fail-safe option on a particular interface card. You can view interface card names in the F5 Configuration utility, or you can use the **bigpipe interface** command to display interface names on the command line.

To arm fail-safe on an interface in the F5 Configuration utility

1. In the navigation pane, click NICs (network interface cards).
   The Network Interface Cards list opens and displays each installed NIC.
2. Select an interface name.
   The Network Interface Card Properties screen opens.
3. Check **Arm Failsafe** to turn on the fail-safe option for the selected interface.
4. In the **Timeout** box, type the maximum time allowed for a loss of network traffic before a fail-over occurs.
5. Click the **Apply** button.

To arm fail-safe on an interface on the command line

One of the required parameters for the **bigpipe interface** command is the name of the interface itself. If you need to look up the names of the installed interface cards, use the **bigpipe interface** command with the **show** keyword:

```
bigpipe interface show
```

To arm fail-safe on a particular interface, use the **bigpipe interface** command with the **failsafe arm** keyword and interface name parameter:

```
bigpipe interface timeout <seconds>
bigpipe interface <ifname> failsafe arm
```

For example, you have an external interface named **exp0** and an internal interface named **exp1**. To arm the fail-safe option on both cards with a timeout of 30 seconds, you need to issue the following commands:

```
bigpipe interface timeout 30
bigpipe interface exp0 failsafe arm
bigpipe interface exp1 failsafe arm
```

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 sendmail 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 **BIG/ip Controller Administrator Guide**, Using per-connection routing, on page 2-30.*

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 really necessary to set the default route for nodes directly to the BIG/ip Controller, so 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.

Note

Configuring GateD on the BIG/ip Controller is not required. Most routing requirements for the BIG/ip Controller can be met without using GateD.

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 F5 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 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 3.1.

```
; 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 3.1 *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 3.2, for an example of how to make the entry in the **/etc/irs.conf** file:

```
hosts           local   continue
hosts           dns
```

Figure 3.2 *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 3.3.

```
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 3.3 *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 3.4.

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

Figure 3.4 *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 3-55.

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 3.5 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 3.5 *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
   ```

Basic configuration examples

Now that you have completed the most basic configuration options for the BIG/ip Controller, you may want to review a couple of sample configurations that use the BIG/ip Controller for traffic management. The following examples are described in this section:

- A basic e-commerce and web server combination
- A simple intranet

The most common application of the BIG/ip Controller is to distribute traffic across an array of web servers that host standard web traffic, including e-commerce traffic. However, a BIG/ip Controller can also control traffic distribution for other types of devices, such as cache servers, proxy servers, firewalls, and even routers.

The following sections provide you with two basic configuration examples that can help you plan your installation. These examples can also help you understand how people use some of the most popular BIG/ip Controller features to resolve specific issues or to enhance network performance in general.

A basic web site and e-commerce configuration

First, we start with a basic configuration where a BIG/ip Controller load balances two sites: **www.MySite.com** and **store.MySite.com**. The **www.MySite.com** site provides standard web content, and the **store.MySite.com** site is the e-commerce site that sells items to **www.MySite.com** customers. In this scenario, the BIG/ip Controller provides simple load balancing for both sites.

Setting up the topology

To set up load balancing for these sites, you need to create two virtual servers, one for each site. Even though the sites are related and they may even share the same IP address, each requires its own virtual server because it uses a different port to support its particular protocol: port 80 for the HTTP traffic going to **www.MySite.com**, and port 443 for the SSL traffic going to **store.MySite.com**.

Figure 3.6 shows the topology for the sample configuration. Each site uses two of the three web servers to host its content. Both sites happen to share Server 2.

◆ Note

Note that in this example, as in all examples in this guide, we use only non-routable IP addresses. In a real topology, the virtual server IP addresses would have to be routable on the Internet.

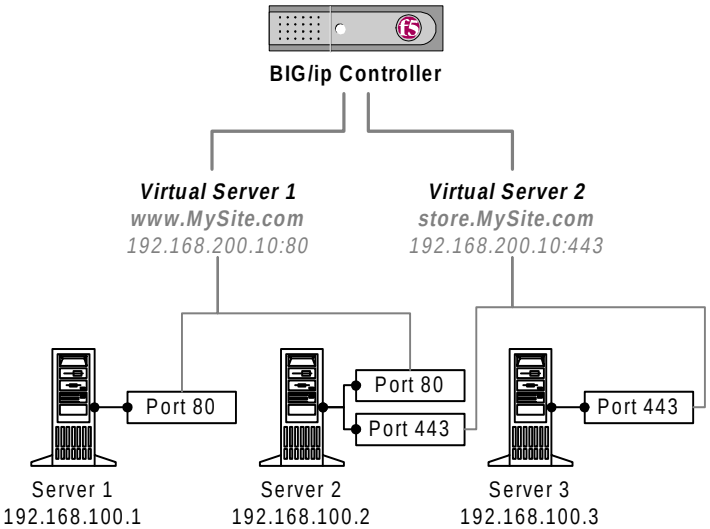

Figure 3.6 *A basic configuration*

The ***virtual servers*** that you define always include three basic elements:

- **virtual IP address**
  This is the IP address that is registered with DNS and associated with your site's domain name. In our example, both **www.MySite.com** and **store.MySite.com** use the same IP address: **192.168.200.10**. Both domain names would presumably have to be registered with DNS to resolve to that IP address.
- **Port**
  The port that hosts the specific service supported by your site. In our example, we have two different sites that support two different ports: port 80 and port 443.
- **Servers that host your site**
  The list of physical servers that actually host your site connections. For a given server, you list each IP address and port number pair, referred to as a ***node***, that the server handles. Even though our example includes only three servers, it actually has four nodes:

- Server 1 hosts only one node: **192.168.100.1:80**.
- Server 2 hosts two nodes, one for each service it supports: **192.168.100.2:80** and **192.168.100.2:443**.
- Server 3 hosts only one node: **192.168.100.3:443**.

The BIG/ip Controller distributes connections among the three servers according to a user-specified load balancing mode. The most common mode is Round Robin, which simply distributes each new connection to the next server in line, eventually distributing the connections equally among all the servers.

Using additional features

In this type of configuration, you might want to take advantage of the following BIG/ip Controller features:

❖ **Extended Content Verification**
Verifies that the web servers are not only up and running, but also able to send valid content to clients. For example, you could use Extended Content Verification to make sure that **www.MySite.com** returns its home page rather than an HTTP 404 error.

❖ **Persistence**
Allows returning e-commerce customers to bypass load balancing and connect to the back-end server to which they originally connected that may contain user-specific information. In our example, **store.MySite.com** provides the ability for users to fill a shopping cart, disconnect from the site, and then return up to 24 hours later to purchase the items. When the user returns to purchase the items, the user may need to go to the same back-end server, depending on how the e-commerce site is set up.

❖ **Network Address Translation**
Allows you to make direct administrative connections to the web servers through the BIG/ip Controller. If your administrative workstation is on the network connected to the BIG/ip Controller's external interface, and administrative workstations frequently are, this feature is essential.

- **Secure Network Address Translation (SNAT)**
  Allows you to make map internally routable IP addresses to an externally routable IP address. SNATs do not allow incoming connections.

A basic intranet configuration

The next example is a configuration that might be found in a large corporate intranet. In this scenario, the BIG/ip Controller performs load balancing for two different types of connection requests:

- **Connections to the company's intranet web site**
  The load balancing for the company's intranet web site is similar to basic Internet web site load balancing. The BIG/ip Controller simply load balances the two web servers that host the company intranet web site.
- **Connections to hosts on the Internet**
  In this example, the BIG/ip Controller provides load balancing for connections bound for the Internet. However, the example shows a somewhat sophisticated setup where the BIG/ip Controller actually intercepts HTTP traffic and directs it to a special cache server. Only clients using protocols other than HTTP, such as FTP or SMTP email, get load balanced to one of the two firewalls that lead to the Internet. This greatly reduces the number of concurrent connections that the firewalls have to maintain. Clients looking to retrieve web content get the content from the cache server itself, instead of the actual web site host. If the cache server does not have the content that the client is looking for, the cache server retrieves the content from the real web site on behalf of the client and then forwards it to the client.

Setting up the topology

To set up load balancing for this intranet example, you need to create three virtual servers: one that handles load balancing for the internal corporate web site, one that directs outbound HTTP traffic to the cache server, and one that handles load balancing for the firewalls.

Figure 3.7 shows the topology for the sample configuration. A standard virtual server handles the load balancing for the corporate intranet web site, **Corporate.main.net**. Wildcard Virtual Server 1 takes all of the outbound HTTP traffic and directs it to the cache server. Wildcard Virtual Server 2 handles all of the remaining traffic that actually has to go out to the Internet.

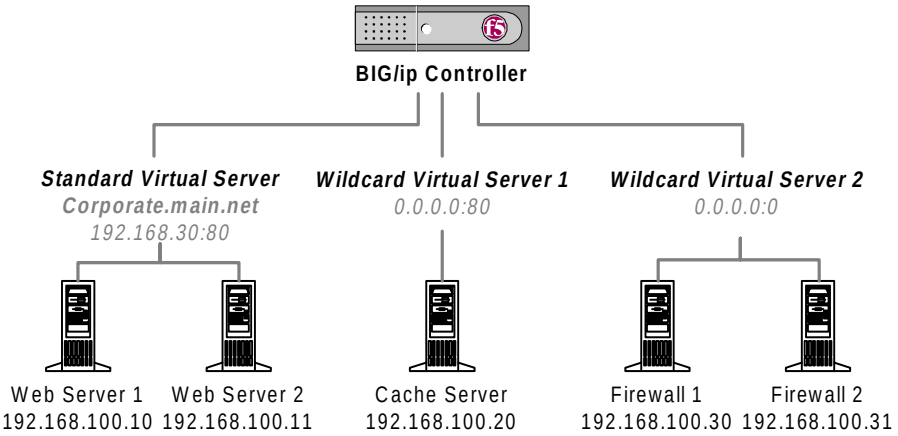

Figure 3.7 *A basic intranet configuration*

The ***wildcard virtual servers*** are a special type of virtual server, which accept traffic going to IP addresses unknown to the BIG/ip Controller, as all outside Internet addresses would be. When the BIG/ip Controller receives a connection request, it immediately tries to match the requested IP address to one of its virtual server IP addresses. If it cannot find a match among the standard virtual servers that it manages, it then looks for a wildcard virtual server. Wildcard virtual servers provide the default IP address of **0.0.0.0** that the BIG/ip Controller can use as a sort of catch-all IP address match.

There are actually two types of wildcard virtual servers, and this example shows both:

- **Port-specific wildcard virtual servers**
  A port-specific wildcard virtual server uses the default IP address, but it has a specific port number, and it only handles

traffic associated with that port number. In our preceding example, the port-specific wildcard virtual server captures all outbound traffic that uses port 80 and directs it to the cache server.

- **Default wildcard virtual servers**
  A default wildcard virtual server is one that uses only port **0**. Port **0**, like the **0.0.0.0** IP address, is a catch-all match for outgoing traffic that does not match any standard virtual server or any port-specific wildcard virtual server. Default wildcard virtual servers typically handle traffic only for firewalls or routers. In the preceding example, the default wildcard virtual server load balances the intranet's firewalls that connect to the Internet.

Using additional features

In this type of configuration, you might want to take advantage additional BIG/ip Controller features that are described in the BIG/ip Controller Administrator Guide and the BIG/ip Controller Reference Guide. These features include:

- **State mirroring**
  This feature is available only for redundant BIG/ip Controller systems, and it greatly enhances the reliability of your network. A redundant system runs two BIG/ip Controllers at the same time. One unit actively handles all connection requests, and the other unit acts as a standby that immediately takes over if the active unit fails and reboots. The state mirroring feature allows the standby unit to maintain all of the current connection and persistence information. If the active unit fails and the standby unit takes over, all connections continue, virtually uninterrupted. This is especially useful for long-lived connections, such as FTP connections which would otherwise have to re-establish an entire transfer session.
- **Destination address affinity**
  Allows the BIG/ip Controller to cache content on specified cache servers by sending all requests for the same server to the same node. This avoids caching the same content on multiple cache servers. Because the above example includes only one cache server, you would not actually implement this feature in that

example. However, the destination address affinity feature is very useful for users who work with multiple cache servers in a similar intranet scenario. Caching specific information on the same cache server saves disk space on your cache servers.

- **IP address filtering**
  Allows you to deny connections going to or coming from specific IP addresses. This feature is useful if you are experiencing denial-of-service attacks from hostile sources. You can set up an IP filter to ignore traffic coming in from the hostile IP address.

Index

/etc/bigd.conf 3-30, 3-34, 3-43
    testing syntax 3-36
/etc/bigip.conf 2-22, 3-43
/etc/ethers 2-23
/etc/hosts 2-23, 3-52
/etc/hosts file 2-24, 2-25
/etc/hosts.allow 3-43
/etc/ipfw.conf 3-43
/etc/ipfwrate.conf 3-43
/etc/irs.conf 3-52
/etc/netstart 2-23
/etc/rateclass.conf 3-43
/etc/snmpd.conf 3-43
/etc/sshd_config 2-23
/etc/ttys 2-10
/var/f5/httpd/conf/httpd.conf 2-23

A

active unit 3-43
Address Resolution Protocol 3-47
Addressing routing issues 3-47
administration
    BIG/ip web server 2-20
administrative connections 2-20
air temperature 2-7
alias 2-16
Allowing access to ports and services 3-12
Arming fail-safe on an interface 3-45
ARP requests 3-45
authoritative name server 3-53

B

Basic configuration examples 3-56
BIG/ip Controller
    default route 2-14
    host name 2-14
BIG/ip HA Controller 1-2
BIG/ip HA+ Controller 1-3
BIG/ip LB Controller 1-2
BIG/ip web server 2-20, 2-26
    configuration 2-20
    configuring settings 2-20
    setting the password 2-20
Broadcast address 2-18

C

cache servers 3-62
command line access 2-25
Committing your configuration settings to the system 2-22
components 2-1
config_failover 3-42
config_ftpd 2-26
configuration
    routing 3-46
configuration files 2-22
Configuration for BIG/ip Controller redundant systems 2-11
configuration synchronization 3-41
configuration synchronization feature 2-20
configurations
    scalability 1-2
Configure SSH screen 2-19
Configuring a time zone 2-14
Configuring an interface for the external network 2-17
Configuring an interface for the internal network 2-18
Configuring and synchronizing redundant systems 3-41
Configuring DNS on the BIG/ip Controller 3-51
Configuring DNS resolution 3-51
Configuring internal and external interfaces 2-15
Configuring SNAT address mappings 3-27
Confirming your configuration settings 2-22
connection request 3-8
connections 1-3, 3-5
content servers 3-48
Converting from rotary or round robin DNS 3-53

D

Data Fellows 2-2, 2-25
default route 2-12, 2-16, 3-46
    configuring a default route 2-14
default router 3-29, 3-45
defining additional host names 2-23
Destination address affinity 3-62
destination IP address 3-8
DNS 3-51
DNS forwarding proxy 2-14
DNS proxy 3-47, 3-53
DNS, converting from rotary DNS 3-53
domain names 3-53

E

Earth ground 2-8
EAV service check 3-17
e-commerce configuration 3-57
ECV 3-5
ECV normal 3-30
ECV service check 3-17
    regular expressions 3-31
ECV service check properties 3-30
ECV SSL 3-31
email
    configuring 3-54
encryption 1-3, 2-2
Ethernet 2-2
Extended Content Verification 3-59
Extended Content Verification (ECV) 3-29
external interface 2-9, 3-22
external network
    configuring an interface 2-17
external network interface 2-17, 2-25

F

F5 Configuration utility 3-33, 3-45
    online help 1-4
fail-over 2-15, 3-45
fail-over cable 2-10
fail-over IP address 2-11, 2-15
fail-safe 3-44
fail-safe settings 3-42
FDDI/CDDI 2-2
file 2-10
First-Time Boot utility 2-1, 2-11, 2-19, 2-23, 3-46
    configuring 2-11
    required information 2-12
    saved files 2-22
    starting 2-12
forwarding virtual server 3-23
FQDN 3-51
F-Secure license agreement 2-25
F-Secure manual 2-26
F-Secure SSH 2-25, 2-26
F-Secure SSH client 1-3, 2-26, 2-28
    downloading 2-26
    for UNIX 2-29
    for Windows 95 2-28
    for Windows NT 2-28
    manual 2-2
FTP 2-26
    allowing on ports 3-13

G

GateD 3-50
gateway 3-48
global load balancing 3-19

H

HA 2-3
HA+ 2-3
hardware
    configuration 2-3
    installing 2-1
    installing and connecting 2-8
    peripheral 2-2
    requirements 2-1
hardware configuration 1-1
    2U controller 2-3, 2-5
    4U controller 2-3, 2-3
Hardware provided with the BIG/ip Controller 2-1

host name 2-11, 2-21, 3-5
defining 2-14
restrictions 2-14
host names
defining additional 2-23

I

idle connection 3-14
idle connection timers
TCP 3-16
intranet configuration 3-61
installation
planning 2-7
procedures 2-8
rack mounting 2-7
interface
arming fail-safe 3-45
attributes 2-15
external 2-15
internal 2-15
interface card 3-45
interface cards 2-2
interfaces
configuring interfaces 2-14
internal interface 2-9, 2-15
configuring 2-18
internal network
configuring an interface 2-18
internal network interface 2-25
international 2-19
intranet configuration 3-60
IP address 2-15, 2-16, 2-17, 2-21, 3-4, 3-52
for fail-over 2-15
IP address filtering 3-63
IP address for remote administration 2-11
IP forwarding 3-22, 3-28

L

LB 2-3
LB Controllers 2-19
LED indicators 2-17
license statement 2-12
lithium battery 2-8
load balancing 3-4, 3-6, 3-8, 3-30, 3-60
load balancing mode 3-1

M

mail exchanger 3-54
manufacturer's instructions 2-7
Media type for Interface 2-17, 2-19
media types
supported 2-19

N

name resolution 3-52
NAT 3-2, 3-5, 3-22, 3-23
netmask 2-16, 2-18
network adapters 2-17
network address translation 3-23
network address translations 3-22
network addresses 3-53
network interface cards (NICs) 2-8, 2-9, 2-23
network traffic 3-44
networking issues 3-46
node addresses 3-23
network address translation 3-23
node ping timer 3-14
nodes 3-22, 3-46
transparent 3-8

O

online help 1-4
outside the US support 2-25

P

Passwords 2-12
passwords
BIG/ip web server 2-20
persistence 3-2, 3-5, 3-36, 3-59
persistent connections 3-37
PKZip 2-28
platform options 1-2
ports 2-4, 2-6, 3-1

power cable 2-10
Preparing to use the synchronization command 3-42
Preparing workstations for command line access 2-25
primary internal interface 2-19

R

rack 2-7
rack mounting 2-7
Ratio mode 3-19, 3-20
ratio weight 3-20
receive expression 3-31
    sample 3-32
reconfig-httpd 2-21
redundant system 2-16, 2-19, 3-2, 3-62
redundant systems 2-2, 2-17, 3-42
    fail-safe interfaces 3-45
release notes 1-3
remote administration 2-2, 2-19
resources
    help 1-3
Root password 2-11
root password 3-42
    defining 2-13
root system 2-12
rotary DNS, converting 3-53
round robin 3-19, 3-53
router configurations 3-47
router's IP address 2-12
routing
    content server to BIG/ip Controller 3-48
    content servers on different logical networks 3-49
    GateD 3-50
routing table 3-29

S

scalability
    BIG/ip Controller 1-2
secure network address translation 3-23
secure network address translation (SNAT) 3-25
security
    BIG/ip web server 2-20
security features 3-12
send string 3-31
    sample 3-32
serial terminal
    configuring 2-10
    in addition to the console 2-10
serial terminal setting 2-10
serial terminals 2-2
servers 2-23
service checking 3-2
service names 3-5
services 3-1
Shared IP alias 2-17
Shared IP alias netmask 2-17
simple persistence 3-37
simple service check 3-17
SNAT 3-23
    address mappings 3-27
    connection limits 3-25
    TCP idle connection timeout 3-26
    UDP idle connection timeout 3-26
SNAT global properties 3-25
software configuration
    basic 1-1
specifications
    BIG/ip Controller 1-1
SSH 1-3
SSH client 2-27, 2-29
    downloading via FTP 2-26
    downloading via the BIG/ip web server 2-26
    installing 2-29
    transferring 2-27
    UNIX 2-29
    untar 2-29
    Windows 95 and Windows NT 2-28
SSL persistence 3-37
standby unit 3-43
State mirroring 3-62

Synchronizing configurations between controllers 3-43

T

TAR/Gzip 2-29
TCP traffic 3-15
technical support 1-3, 1-4
Third-party documentation 1-4
time zone 2-14
    configuring a time zone 2-14
timer 3-13
    node ping timer 3-14
    reaping idle connections 3-15
    service check 3-17
timer settings 3-1
transparent network devices 3-4, 3-8
transparent nodes
    defining 3-9

U

UDP traffic 3-15
unit ID 2-15
US 2-19
US support 2-25
Using standard or wildcard virtual servers 3-4
utilities
    First-Time Boot utility 2-11

V

VGA monitor 2-9
virtual server 2-25, 3-1, 3-38
    configuring 3-4–3-12
    defining standard 3-6–3-8
    defining wildcard virtual servers 3-8
    overview 3-58
    using additional features 3-5
    wildcard 3-12
    wildcard or standard 3-4
virtual server IP address 3-4
virtual server mapping
    defined 3-6
virtual server mappings 3-1, 3-10
voltage selector 2-8

W

web access 2-21
web servers 3-30
wildcard ports 3-9
WinZip 2-28
World Wide Web 1-4

Z

ZIP tools 2-28
zone tables 3-53