BIG-IP® e-Commerce Solutions Guide

version 4.2

MAN-0060-00

Service and Support Information

Product Version

This manual applies to version 4.2 of the BIG-IP® e-Commerce Controller.

Obtaining Technical Support

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

Contacting F5 Networks

Web	www.f5.com
Toll-free phone	(888) 88BIGIP
Corporate phone	(206) 272-5555
Fax	(206) 272-5556
Email	sales@f5.com
Mailing Address	401 Elliott Avenue West Seattle, Washington 98119

Legal Notices

Copyright

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

Trademarks

F5 and the F5 logo, F5 Networks, BIG-IP, 3-DNS, GLOBAL-SITE, EDGE-FX, and SEE-IT are registered trademarks of F5 Networks, Inc. FireGuard, iControl, Internet Control Architecture, and IP Application Switch are trademarks of F5 Networks, Inc. In Japan, the F5 logo is trademark number 4386949, BIG-IP is trademark number 4435184, 3-DNS is trademark number 4435185, and SEE-IT is trademark number 4394516. All other product and company names are registered trademarks or trademarks of their respective holders.

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export a this product from the United States.

Export Warning

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

FCC Compliance

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

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

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

Canadian Regulatory Compliance

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

Standards Compliance

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

Acknowledgments

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

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

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

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

This product includes software developed by Adam Glass.

This product includes software developed by Christian E. Hopps.

This product includes software developed by Dean Huxley.

This product includes software developed by John Kohl.

This product includes software developed by Paul Kranenburg.

This product includes software developed by Terrence R. Lambert.

This product includes software developed by Philip A. Nelson.

This product includes software developed by Herb Peyerl.

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

This product includes software developed by Chris Provenzano.

This product includes software developed by Theo de Raadt.

This product includes software developed by David Muir Sharnoff.

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

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

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

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

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

This product includes software developed by Christos Zoulas.

This product includes software developed by Charles Hannum.

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

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

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

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

This product includes software developed by the Apache Group for use in the Apache HTTP server project (http://www.apache.org/).

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

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

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

Table of Contents

Introduction

1
Configuring an SSL Accelerator

Glossary

Index

Introduction

- Getting started
- Using the Administrator Kit
- What's new in version 4.2
- Learning more about the BIG-IP product family

Getting started

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

Once you find your solution and gather the necessary network information, turn back to the ***Configuration Worksheet*** and ***Hardware Orientation*** poster for hardware installation instructions, and then return to the ***BIG-IP e-Commerce Solutions Guide*** to follow the steps for setting up your chosen solution.

Choosing a configuration tool

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

The Setup utility

All users will use the Setup utility (formerly known as First-Time Boot utility). This utility walks you through the initial system set up. You can run the Setup utility from the command line, or from a web browser. The Setup utility prompts you to enter basic system information including a root password and the IP addresses that will be assigned to the network interfaces. For more information, see Chapter 2 of the ***BIG-IP Reference Guide***.

The Configuration utility

The Configuration utility is a web-based application that you use to configure and monitor the setup on the BIG-IP. Once you complete the installation instructions described in this guide, you can use the Configuration utility to perform the configuration steps necessary for your chosen load balancing solution. In the Configuration utility, you can also monitor current system performance, and download administrative tools such as the SNMP MIB or the SSH client. The Configuration utility requires Netscape Navigator version 4.7, or Microsoft Internet Explorer version 5.0 or later.

The bigpipe and bigtop command line utilities

The **bigpipe**™ utility is the command line counter-part to the Configuration utility. Using **bigpipe** commands, you can configure virtual servers, open ports to network traffic, and configure a wide variety of features. To monitor the BIG-IP, you can use certain **bigpipe** commands, or you can use the **bigtop**™ utility, which provides real-time system monitoring. You can use the command line utilities directly on the BIG-IP console, or you can run commands using a remote shell, such as the SSH client (encrypted

communications only), or a Telnet client (for countries restricted by cryptography export laws). For detailed information about the command line syntax, see the Chapter 7 of the ***BIG-IP Reference Guide***.

Using the Administrator Kit

The BIG-IP Administrator Kit provides all of the documentation you need to work with the BIG-IP. The information is organized into the guides described below. The following printed documentation is included with the BIG-IP unit.

- **Hardware Orientation Poster**
  This poster includes information about the BIG-IP unit. It also contains important environmental warnings.
- **Configuration Worksheet**
  This worksheet provides you with a place to plan the basic configuration for the BIG-IP.

The following guides are available in PDF format from the CD-ROM provided with the BIG-IP. These guides are also available from the first Web page you see when you log in to the administrative web server on the BIG-IP.

- **BIG-IP e-Commerce Solutions Guide**
  This guide provides examples of common load balancing solutions. Before you begin installing the hardware, we recommend that you browse this guide to find the load balancing solution that works best for you.
- **BIG-IP Reference Guide**
  This guide provides detailed configuration information for the BIG-IP. It also provides syntax information for **bigpipe** commands, other command line utilities, configuration files, system utilities, and monitoring and administration information.

Stylistic conventions

To help you easily identify and understand important information, our documentation uses the stylistic conventions described below.

Using the solution examples

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

Identifying new terms

To help you identify sections where a term is defined, the term itself is shown in bold italic text. For example, a ***virtual server*** is a specific combination of a virtual address and virtual port, associated with a content site that is managed by a BIG-IP or other type of host server.

Identifying references to objects, names, and commands

We apply bold text to a variety of items to help you easily pick them out of a block of text. These items include web addresses, IP addresses, utility names, and portions of commands, such as variables and keywords. For example, with the **bigpipe pool <pool_name> show** command, you can specify a specific pool to show by specifying a pool name for the **<pool_name>** variable.

Identifying references to other documents

We use italic text to denote a reference to another document. In references where we provide the name of a book as well as a specific chapter or section in the book, we show the book name in bold, italic text, and the chapter/section name in italic text to help quickly differentiate the two. For example, you can find information about **bigpipe** commands in the ***BIG-IP Reference Guide***, Chapter 7, *bigpipe Command Reference.*

Identifying command syntax

We show complete commands in bold Courier text. Note that we do not include the corresponding screen prompt, unless the command is shown in a figure that depicts an entire command line screen. For example, the following command shows the configuration of the specified pool name:

```
bigpipe pool <pool_name> show
```

or

```
b pool <pool_name> show
```

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

Item in text	Description
****	Indicates that the command continues on the following line, and that users should type the entire command without typing a line break.
< >	Identifies a user-defined parameter. For example, if the command has **<your name>**, type in your name, but do not include the brackets.
|	Separates parts of a command.
[ ]	Indicates that syntax inside the brackets is optional.
...	Indicates that you can type a series of items.

Table Intro.1 *Command line syntax conventions*

Finding additional help and technical support resources

You can find additional technical information about this product in the following locations:

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

What's new in version 4.2

The BIG-IP offers the following major new features in version 4.2, in addition to many smaller enhancements.

Improved BIG-IP licensing procedure

With this release, you can now re-license your BIG-IP systems using the web-based Configuration utility. This new method of licensing simplifies the BIG-IP licensing process in general.

The Setup Utility

This release includes a new Setup utility for initially configuring your BIG-IP system. The Setup utility replaces the web-based and console-based First-Time Boot utility. For more information, see the ***BIG-IP Reference Guide***, Chapter 2, *Using the Setup Utility*.

Enhanced pools support

This release provides a feature to specify a pool of multiple default gateways, used for handling administrative traffic such as SSH, telnet, FTP, and HTTPS connections. For more information, see the *Pools* section in the ***BIG-IP Reference Guide***, Chapter 4, *Configuring the High-Level Network.*

SSL Accelerator proxy enhancements

This release includes several important enhancements to the SSL Accelerator proxy. For example, you can now configure options such as specifying ways for an SSL proxy to manage client certificates, inserting headers into HTTP requests, specifying ciphers and protocol versions, and configuring SSL session cache size and timeout values.

This release also supports the SSL-to-Server option, which allows you to re-encrypt traffic after it has been decrypted by the BIG-IP. Previously available on the IP Application Switch only, this feature is now available on the e-Commerce and BIG-IP Controller platform also. Moreover, this feature has been enhanced in this release to further ensure the security of SSL connections between the proxy and the server.

For a complete list of all new SSL Accelerator proxy options, see the *Proxies* section in Chapter 4 of the ***BIG-IP Reference Guide***.

New filter for rewriting HTTP redirections

This release provides an ISAPI filter, called **redirectfilter.dll**, which allows IIS servers running Netscape to rewrite HTTP redirections. Rewriting HTTP redirections helps to ensure that SSL connections remain on a secure channel. By installing this filter on your IIS server, you offload the task of rewriting HTTP redirections from your SSL Accelerator proxy to your IIS server. For more information, see the *Rewriting HTTP redirection* section (which is a subsection of the *HTTP Redirection* section) in Chapter 4 of the ***BIG-IP Reference Guide***.

Support for the NCipher FIPS 140-compatible hardware security module

For BIG-IP Controller platforms, option is available to install a FIPS 140-1-certified cryptographic network module. The BIG-IP FIPS hardware option is specifically designed for processing SSL traffic within

environments that require FIPS 140-1 Level 3 compliant solutions. It comes with the FIPS 140-1 level 3 certified PCI based encryption processing module, attached smart card reader, and 5 smart cards. This product can be installed in any BIG-IP Controller platform that has BIG-IP software version 4.2 and is authorized by your vendor. For more information, see *Configuring FIPS 140 Security World on the BIG-IP* in the Documentation section of the Software and Documentation CD.

Enhanced support for Secure Network Address Translations (SNATs)

In previous releases, BIG-IP allowed you to automatically map VLANs to translation IP addresses during SNAT creation. In this release, you can now use this automapping feature not only for VLANs, but for one or more individual IP addresses. For more information, see the *Address Translation: NATs, SNATs, and IP Forwarding* section in the ***BIG-IP Reference Guide***, Chapter 4, *Configuring the High-Level Network*.

Enhanced interface statistics

This release features enhanced statistics for BIG-IP interfaces. The following state information and statistics are now available: MTU, Speed, MAC address, packets in, errors in, packets out, errors out, collisions, dropped packets, bits in, bits out. Previously available on the IP Application Switch, this feature is new for the BIG-IP Controller platform. For more information, see the ***BIG-IP Reference Guide***, Chapter 11, *Monitoring and Administration*.

Support for LDAP and RADIUS logins

With this release, BIG-IP can now authenticate SSH users by way of an LDAP or a RADIUS server. For information on configuring this feature, see the *To configure RADIUS login support* section and the *Configuring LDAP login support* section (which are a subsections of the *Configuring RADIUS or LDAP authentication* section) in Chapter 12 of the ***BIG-IP Reference Guide***.

Enhanced system logging

System logging in this release provides more detailed information, such as **up** or **down** status for nodes. For more information, see the BIG-IP Reference Guide, Chapter 11, *Monitoring and Administration.*

Web-based Configuration utility enhancements

This release includes a number of improvements to the web-based Configuration utility. All new features for this release are supported by the Configuration utility.

Learning more about the BIG-IP product family

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

- **The BIG-IP**
  The BIG-IP HA, and HA+ software provides the full suite of local area load balancing functionality. The BIG-IP unit also has an optional 3-DNS software module which supports wide-area load balancing.
- **The BIG-IP e-Commerce Controller**
  The BIG-IP e-Commerce Controller uses SSL acceleration technology to increase the speed and reliability of the secure connections that drive e-commerce sites.
- **The BIG-IP special purpose products**
  The special purpose BIG-IP provides the ability to choose from three different BIG-IP feature sets. When you run the Setup utility, you specify one of three types:
  - **The BIG-IP Load Balancer**
    The BIG-IP Load Balancer provides basic load balancing features.
  - **The BIG-IP FireGuard**
    The BIG-IP FireGuard provides load balancing features that maximize the efficiency and performance of a group of firewalls.
  - **The BIG-IP Cache Controller**
    The BIG-IP Cache Controller uses content-aware traffic direction to maximize the efficiency and performance of a group of cache servers.

1

Configuring an SSL Accelerator

- Introducing the SSL Accelerator
- Configuring the SSL Accelerator
- Introducing the SSL Accelerator scalable configuration
- Using SSL-to-server

Introducing the SSL Accelerator

The SSL Accelerator feature allows the BIG-IP e-Commerce Controller to accept HTTPS connections (HTTP over SSL), connect to a web server, retrieve the page, and then send the page to the client.

A key component of the SSL Accelerator feature is that the e-Commerce Controller can retrieve the web page using an unencrypted HTTP request to the content server. With the SSL Accelerator feature, you can configure an SSL gateway on the e-Commerce Controller that decrypts HTTP requests that are encrypted with SSL. Decrypting the request offloads SSL processing from the servers to the BIG-IP and also allows the BIG-IP to use the header of the HTTP request to intelligently control how the request is handled. (Requests to the servers can optionally be re-encrypted to maintain security on the server side of the e-Commerce Controller as well, using a feature called SSL-to-server.)

When the SSL proxy on the BIG-IP connects to the content server, and address translation is not enabled, the proxy uses the original client's IP address and port as its source address and port. In doing so, the proxy appears to be the client, for logging purposes.

This chapter describes the following features of the BIG-IP SSL Accelerator:

- Configuring an SSL Accelerator
- Using an SSL Accelerator scalable configuration
- Using SSL-to-server

◆ Note

If you have FIPS-140 security modules installed in the e-Commerce Controller, you must initialize the security world before you configure the SSL Accelerator for encrypted traffic. For more information, see ***Configuring the FIPS-140 Security World****, available on the Software and Documentation CD.*

◆ Note

All products except the BIG-IP LoadBalancer, BIG-IP FireGuard Controller, and the BIG-IP Cache Controller support this configuration.

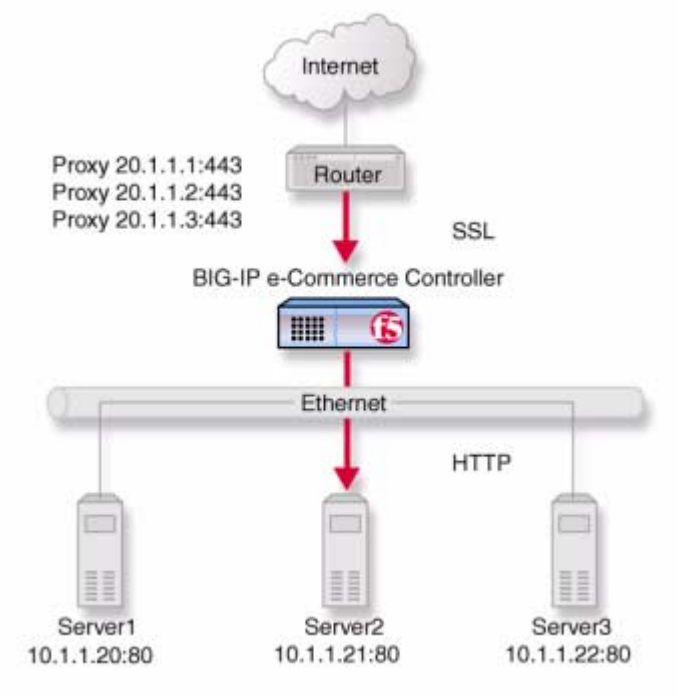

Figure 1.1 *An incoming SSL connection received by an SSL Accelerator configured on the BIG-IP*

Configuring the SSL Accelerator

There are several tasks required to set up the SSL Accelerator on the BIG-IP. These tasks include:

- Generating a key and obtaining a certificate
- Configuring the BIG-IP with the certificate and key
- Creating the gateway for the SSL Accelerator

Generating a key and obtaining a certificate

In order to use the SSL Accelerator feature, you must obtain a valid x509 certificate from an authorized certificate authority (CA).

◆ Note

If you have FIPS-140 hardware installed in the BIG-IP, see ***Configuring the Security World on BIG-IP*** *on the product and documentation CD for instructions on how to generate a key and obtain a certificate.*

The following list contains some companies that are certificate authorities:

- Verisign (**http://www.verisign.com**)
- Digital Signature Trust Company (**http://secure.digsigtrust.com**)
- GlobalSign (**http://www.globalsign.com**)
- GTE Cybertrust (**http://www.cybertrust.gte.com**)
- Entrust (**http://www.entrust.net**)

You can generate a key, a temporary certificate, and a certificate request form using the Configuration utility or from the command line.

Note that we recommend using the Configuration utility for this process. The certification process is generally handled through a web page. Parts of the process require you to cut and paste information from a browser window in the Configuration utility to another browser window on the web site of the CA.

Additional information about keys and certificates

You must have a separate certificate for each domain name on each BIG-IP or redundant pair of BIG-IP units, regardless of how many non-SSL web servers are load balanced by the BIG-IP.

If you are already running an SSL server, you can use your existing keys to generate temporary certificates and request files. However, you must obtain new certificates if the ones you have are not for the following web server types:

- Apache + OpenSSL
- Stronghold

Generating a key and obtaining a certificate using the Configuration utility

To obtain a valid certificate, you must have a private key. If you do not have a key, you can use the Configuration utility on the BIG-IP to generate a key and a temporary certificate. You can also use the Configuration utility to create a request file you can submit to a certificate authority (CA). You must complete three tasks in the Configuration utility to create a key and generate a certificate request.

- Generate a certificate request
- Submit the certificate request to a CA and generate a temporary certificate
- Install the SSL certificate from the CA

Each of these tasks is described in detail in the following paragraphs.

To create a new certificate request using the Configuration utility

1. In the navigation pane, click **Proxies**.
   The Proxies screen opens.

2. Click the Create SSL Certificate Request tab.
   The New SSL Certificate Request screen opens.
3. In the Key Information section, select a key length and key file name.
   - **Key Length**
     Select the key length you want to use for the key. You can choose **512**, **1024, 2048** or **4096** bits.
   - **Keyfile Name**
     Type in the name of the key file. This should be the fully qualified domain name of the server for which you want to request a certificate. You must add the **.key** file extension to the name.
4. In the Certificate Information section, type the information specific to your company. This information includes:
   - **Country**
     Type the two letter ISO code for your country, or select it from the list. For example, the two-letter code for the United States is **US**.
   - **State or Province**
     Type the full name of your state or province, or select it from the list. You must enter a state or province.
   - **Locality**
     Type the city or town name.
   - **Organization**
     Type the name of your organization.
   - **Organizational Unit**
     Type the division name or organizational unit.
   - **Domain Name**
     Type the name of the domain upon which the server is installed.
   - **Email Address**
     Type the email address of a person who can be contacted about this certificate.
   - **Challenge Password**
     Type the password you want to use as the challenge password for this certificate. The CA uses the challenge password to verify any changes you make to the certificate at a later date.
   - **Retype Password**
     Retype the password you entered for the challenge password.
5. Click the **Generate Certificate Request** button.
   After a short pause, the SSL Certificate Request screen opens.
6. Use the SSL Certificate Request screen to start the process of obtaining a certificate from a CA, and then to generate and install a temporary certificate.

- **Begin the process for obtaining a certificate from CA**
  Click the URL of a CA to begin the process of obtaining a certificate for the server. After you select a CA, follow the directions on their web site to submit the certificate request. After your certificate request is approved, and you receive a certificate back from the CA, see *To install certificates from the CA using the Configuration utility*, on page 1-7, for information about installing it on the BIG-IP.
- **Generate and install a temporary certificate**
  Click the **Generate Self-Signed Certificate** button to create a self-signed certificate for the server. We recommend that you use the temporary certificate for testing only. You should take your site live only after you receive a properly-signed certificate from a certificate authority. When you click this button, a temporary certificate is created and installed on the BIG-IP. This certificate is valid for 10 years. This temporary certificate allows you to set up an SSL gateway for the SSL Accelerator while you wait for a CA to return a permanent certificate.

Generating a key and obtaining a certificate from the command line

To obtain a valid certificate, you must have a private key. If you do not have a key, you can use the **genconf** and **genkey** utilities on the BIG-IP to generate a key and a temporary certificate. The **genkey** and **gencert** utilities automatically generate a request file that you can submit to a certificate authority (CA). If you have a key, you can use the **gencert** utility to generate a temporary certificate and request file.

These utilities are described in the following list:

- **genconf**
  This utility creates a key configuration file that contains specific information about your organization. The **genkey** utility uses this information to generate a certificate.
- **genkey**
  After you run the **genconf** utility, run this utility to generate a temporary 10-year certificate for testing the SSL Accelerator on the BIG-IP. This utility also creates a request file that you can submit to a certificate authority (CA) to obtain a certificate.
- **gencert**
  If you already have a key, run this utility to generate a temporary certificate and request file for the SSL Accelerator.

To generate a key configuration file using the genconf utility

If you do not have a key, you can generate a key and certificate with the **genconf** and **genkey** utilities. First, run the **genconf** utility with the following commands:

```
/usr/local/bin/genconf
```

The utility prompts you for information about the organization for which you are requesting certification. This information includes:

- The fully qualified domain name (FQDN) of the server. Note that this FQDN must be RFC1034/1035-compliant, and cannot be more than 63 characters long (this is an x509 limitation).
- The two-letter ISO code for your country
- The full name of your state or province
- The city or town name
- The name of your organization
- The division name or organizational unit

For example, Figure 1.2 contains entries for the server **my.server.net**.

```
Common Name (full qualified domain name): my.server.net
Country Name (ISO 2 letter code): US
State or Province Name (full name): WASHINGTON
Locality Name (city, town, etc.): SEATTLE
Organization Name (company): MY COMPANY
Organizational Unit Name (division): WEB UNIT
```

Figure 1.2 *Example entries for the* ***genconf*** *utility*

To generate a key using the genkey utility

After you run the **genconf** utility, you can generate a key with the **genkey** utility. Type the following command to run the **genkey** utility:

`/usr/local/bin/genkey <server_name>`

For the <**server_name**>, type the FQDN of the server to which the certificate applies. After the utility starts, it prompts you to verify the information created by the **genconf** utility. After you run this utility, a certificate request form is created in the following directory:

`/config/bigconfig/ssl.csr/<fqdn>.csr`

The <**fqdn**> is the fully qualified domain name of the server. Please contact your certificate authority (CA) and follow their instructions for submitting this request form.

In addition to creating a request form that you can submit to a certificate authority, this utility also generates a temporary certificate. The temporary certificate is located in:

`/config/bigconfig/ssl.crt/<fqdn>.crt`

The <**fqdn**> is the fully qualified domain name of the server.

Note that the keys and certificates are copied to the other BIG-IP in a redundant system when you synchronize the configurations. For more information about synchronizing configurations, see the ***BIG-IP Reference Guide***, Chapter 6, *Configuring a Redundant System*.

This temporary certificate is good for ten years, but for an SSL proxy you should have a valid certificate from your CA.

WARNING

Be sure to keep your previous key if you are still undergoing certification. The certificate you receive is valid only with the key that originally generated the request.

To generate a certificate with an existing key using the gencert utility

To generate a temporary certificate and request file to submit to the certificate authority with the **gencert** utility, you must first copy an existing key for a server into the following directory on the BIG-IP:

```
/config/bigconfig/ssl.key/
```

After you copy the key into this directory, type the following command at the command line:

```
/user/local/bin/gencert <server_name>
```

For the **<server_name>**, type the FQDN of the server to which the certificate applies. After the utility starts, it prompts you for various information. After you run this utility, a certificate request form is created in the following directory:

```
/config/bigconfig/ssl.crt/<fqdn>.csr
```

The **<fqdn>** is the fully qualified domain name of the server. Please contact your certificate authority (CA) and follow their instructions for submitting this request form.

Installing certificates from the certificate authority (CA)

After you obtain a valid x509 certificate from a certificate authority (CA) for the SSL Accelerator, you must copy it onto each BIG-IP in the redundant configuration. You can configure the accelerator with certificates using the Configuration utility or from the command line.

To install certificates from the CA using the Configuration utility

1. In the navigation pane, click **Proxies**.
   The Proxies screen opens.
2. On Proxies screen, click the Install SSL Certificate Request tab.
   The Install SSL Certificate screen opens.
3. In the **Certfile Name** box, type the fully qualified domain name of the server with the file extension **.crt**. If you generated a temporary certificate when you submitted a request to the CA, you can select the name of the certificate from the list. This allows you to overwrite the temporary certificate with the certificate from the CA.

4. Paste the text of the certificate into the Install SSL Certificate window. Make sure you include the **BEGIN CERTIFICATE** line and the **END CERTIFICATE** line. For an example of a certificate, see Figure 1.3.
5. Click the **Write Certificate File** button to install the certificate.

```
-----BEGIN CERTIFICATE-----
MIIB1DCCAX4CAQAwDQYJKoZIhvcNAQEEBQAwdTELMAkGA1UEBhMCVVMxCzAJBgNV
BAgTAldBMRAwDgYDVQQHEwdTZWF0dGxlMRQwEgYDVQQKEwtGNSBOZXR3b3JrczEc
MBoGA1UECxMTUHJvZHVjdCBEZXZlbG9wbWVudDETMBEGA1UEAxMKc2VydmVyLm5l
dDAeFw0wMDA0MTkxNjMxNTlaFw0wMDA1MTkxNjMxNTlaMHUxCzAJBgNVBAYTAlVT
MQswCQYDVQQIEwJXQTEQMA4GA1UEBxMHU2VhdHRsZTEUMBIGA1UEChMLRjUgTmV0
d29ya3MxHDAaBgNVBAsTE1Byb2R1Y3QgRGV2ZWxvcG1lbnQxEzARBgNVBAMTCnNl
cnZlci5uZXQwXDANBgkqhkiG9w0BAQEFAANLADBIAkEAsfCFXq3Jt+FevxUqBZ9T
Z7nHx9uaF5x9V5xMZYgekjc+LrF/yazhmq4PCxrws3gvJmgpTsh50YJrhJgfs2bE
gwIDAQABMA0GCSqGSIb3DQEBBAUAA0EAd1q6+u/aMaM2qdo7EjWx14TYQQGomYoq
eydlzb/3FOiJAynDXnGnSt+CVvyRXtvmG7V8xJamzkyEpZd4iLacLQ==
-----END CERTIFICATE-----
```

Figure 1.3 *An example of a certificate*

After the certificate is installed, you can continue with the next step in creating an SSL gateway for the server.

To install certificates from the CA from the command line

Copy the certificate into the following directory on each BIG-IP in a redundant system:

`/config/bigconfig/ssl.crt/`

Note

The certificate you receive from the certificate authority (CA) should overwrite the temporary certificate generated by ***genkey*** *or* ***gencert****.*

If you used the **genkey** or **gencert** utilities to generate the request file, a copy of the corresponding key should already be in the following directory on the BIG-IP:

`/config/bigconfig/ssl.key/`

WARNING

In a redundant system, the keys and certificates must be in place on both BIG-IP units before you configure the SSL Accelerator. To do this you must synchronize the configurations in the redundant system, see the ***BIG-IP Reference Guide***, Chapter 6, *Configuring a Redundant System.*

Creating an SSL gateway

After you create the HTTP virtual server for which the SSL Accelerator handles connections, the next step is to create an SSL gateway. This section also contains information about managing an SSL gateway.

To create an SSL gateway using the Configuration utility

1. In the navigation pane, click **Proxies**.
   The Proxies screen opens.
2. Click the **ADD** button.
   The Add Proxy screen opens.
3. In the Add Proxy screen, configure the attributes you want to use with the proxy. For additional information about configuring a proxy, click the **Help** button.

To create an SSL gateway from the command line

Use the following command syntax to create an SSL gateway:

```
b proxy <ip>:<service> \
target <server | virtual> <ip>:<service> \
clientssl enable \
clientssl key <clientssl_key> \
clientssl cert <clientssl_cert>
```

For example, you can create an SSL gateway from the command line that looks like this:

```
b proxy 10.1.1.1:443 \
target virtual 20.1.1.1:80 \
clientssl enable \
clientssl key my.server.net.key \
clientssl cert my.server.net.crt
```

Introducing the SSL Accelerator scalable configuration

This section explains how to set up at scalable one-armed SSL Accelerator configuration. This configuration is useful for any enterprise that handles a large amount of encrypted traffic.

With this configuration, you can easily add BIG-IP e-Commerce Controllers to keep up with expanding SSL content, or a growing array of SSL content servers without adding more BIG-IP units.

Figure 1.4 shows a scalable configuration. The configuration includes a BIG-IP; the BIG-IP e-Commerce Controllers **Accelerator1**, **Accelerator2**, **Accelerator3**, and **Accelerator4**; and the server array **Server1, Server2, Server3,** and **Server4**.

The following sections refer to Figure 1.4 as an example of how you can set up such a configuration.

Note

The IP addresses shown in these configurations are examples only. When implementing your configuration, choose IP addresses that are consistent with your network or networks.

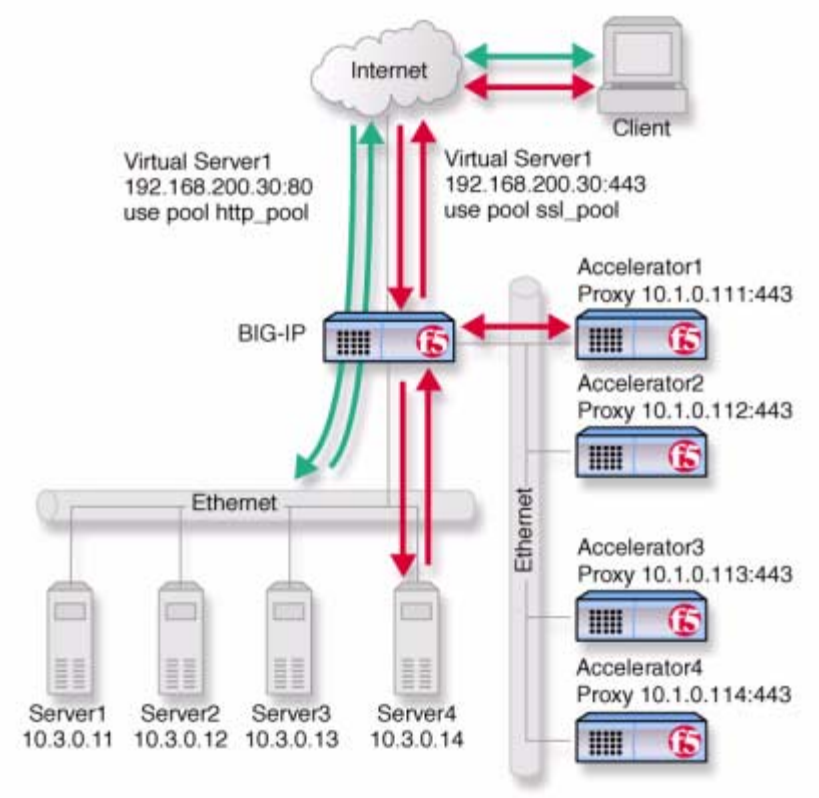

Figure 1.4 *An SSL Accelerator scalable configuration*

Creating the scalable SSL Accelerator configuration

To implement the scalable configuration, you must configure the BIG-IP that load balances the servers and SSL Accelerators, each SSL Accelerator, and each node that handles connections from the SSL Accelerator.

Note

This configuration requires additional load balancing functionality that is not provided with the e-Commerce Controller. This example uses a load balancing version of the BIG-IP.

First, complete the following tasks on the BIG-IP that you want to use to load balance connections to the SSL Accelerators:

- **Create two load balancing pools**
  One pool load balances HTTP connections using the IP addresses of the web servers, the other pool load balances SSL connections to the SSL Accelerators.
- **Create virtual servers**
  Create virtual servers that reference the load balancing pools. Create one virtual server for the pool load balancing the SSL connections to the accelerators, and another virtual server for the pool that load balances the HTTP connections to the servers. Disable external VLAN for the HTTP virtual server to prevent clients from making a direct connection, bypassing the SSL accelerators.
- **Enable service 80 and service 443**
  Enable service **80** and service **443** on the BIG-IP.
- **Set the idle connection timer**
  Set the idle connection timer for service **443**.

Next, complete the following tasks for the SSL Accelerators:

- **Set up SSL gateways**
  Set up an SSL gateway for each accelerator
- **Enable service 443**
  Enable service **443** for encrypted traffic.

Configuring the BIG-IP that load balances the SSL Accelerators

To configure the BIG-IP that load balances the SSL Accelerators, complete the following tasks on the BIG-IP. This section describes how to complete each task.

- Create two load balancing pools. One pool load balances HTTP connections using the IP addresses of the web servers, the other pool load balances SSL connections from the SSL Accelerators.
- Create virtual servers that reference the load balancing pools.
- Enable port 80 and port 443 on the BIG-IP.

Creating load balancing pools

You need to create two pools, a pool to load balance connections using the IP addresses of the content server nodes and a pool to load balance the SSL gateways.

To create the pools using the Configuration utility

1. In the navigation pane, click **Pools**.
   The Pools screen opens.

2. Click the **Add** button.
   The Add Pool screen opens.
3. For each pool, enter the pool name and member addresses in the Add Pool screen. (For additional information about configuring a pool, click the **Help** button.)

   Configuration notes

   For this example, create an HTTP pool named ***http_virtual****. This pool contains the following members:*
   Server1 (10.3.0.11)
   Server2 (10.3.0.12)
   Server3 (10.3.0.13)
   Server4 (10.3.0.14)

   For this example, you could create an SSL accelerator pool named ***ssl_gateways****. This pool contains the following members:*
   accelerator1 (10.1.0.111)
   accelerator2 (10.1.0.112)
   accelerator3 (10.1.0.113)
   accelerator4 (10.1.0.114)

To define a pool from the command line

To define a pool from the command line, use the following syntax:

```
b pool <pool_name> { \
member <member_definition> ...
member <member_definition>}
```

For example, if you want to create the pool **http_virtual** and the pool **ssl_gateways**, you would type the following commands:

```
b pool http_virtual { \
member 10.3.0.11:80 \
member 10.3.0.12:80 \
member 10.3.0.13:80 \
member 10.3.0.14:80 }

b pool ssl_gateways { \
member 10.1.0.111:443 \
member 10.1.0.112:443 \
member 10.1.0.113:443 \
member 10.1.0.114:443 }
```

Creating the virtual servers

Create a virtual server that references the pool that is load balancing the SSL connections, and another virtual server that references the pool that load balances the HTTP connections through the SSL Accelerators.

To define a standard virtual server that references a pool using the Configuration utility

1. In the navigation pane, click **Virtual Servers**.
2. Click the **Add** button.
   The Add Virtual Server screen opens.
3. For each virtual server, enter the virtual server address and pool name. (For additional information about configuring a virtual server, click the **Help** button.)

   Configuration notes

   *To create the configuration described in Figure 1.4, create a virtual server **192.168.200.30** on port **443** that references the pool of SSL accelerators (**ssl_gateways**).*

   *To create the configuration described in Figure 1.4, create a virtual server **192.168.200.30** on port **80** that references the pool of content servers (**http_virtual**).*

To define the virtual servers from the command line

To define a standard virtual server from the command line, use the following syntax:

```
b virtual <virt_IP>:<service> use pool <pool_name>
```

Note that you can use host names in place of IP addresses, and that you can use standard service names in place of port numbers.

To create the virtual servers for the configuration in Figure 1.4, you would type the following commands:

```
b virtual 192.168.200.30:443 use pool ssl_gateways
b virtual 192.168.200.30:80 use pool http_virtual \
vlans external disable
```

Enabling ports 80 and 443 on the BIG-IP

For security reasons, the BIG-IP ports do not accept traffic until you enable them. In this configuration, the BIG-IP accepts traffic on port **443** for SSL, and on port **80** for HTTP. For this configuration to work, you must enable port **80** and port **443**.

Use the following command to enable these ports:

```
b service 80 443 tcp enable
```

Setting the idle connection timer for port 443

In this configuration, you should set the idle connection timer to clean up closed connections on port **443**. You need to set an appropriate idle connection time-out value so that valid connections are not disconnected, and closed connections are cleaned up in a reasonable time.

To set the idle connection time-out using the Configuration utility

1. In the navigation pane, click **Virtual Servers**.
2. In the Virtual Servers list, click the virtual server you configured for SSL connections.
   The Virtual Server Properties screen opens.
3. Click the **Virtual Service Properties** tab.
   The Virtual Service Properties screen opens.
4. In the **Idle connection timeout TCP (seconds)** box, type a time-out value for TCP connections. The recommended time-out setting is 10 seconds.
5. Click **Apply**.

To set the idle connection time-out from the command line

To set the idle connection time-out, type the following command:

```
b service <service> timeout tcp <timeout>
```

The **<timeout>** value is the number of seconds a connection is allowed to remain idle before it is terminated. The **<service>** value is the port on the wildcard virtual server for which you are configuring out-of-path routing. The recommended value for TCP connection timeouts is 10 seconds.

Configuring the SSL Accelerators

The next step in the process is to configure the SSL Accelerators. Complete the following tasks on each SSL Accelerator:

- Set up an SSL gateway for each e-Commerce Controller
- Enable port **443**
- Set the idle connection timer for port **443**

Setting up an SSL gateway for each e-Commerce Controller

The first task you must complete on the SSL Accelerator it to set up an SSL gateway for each e-Commerce Controller with the HTTP virtual server as target server.

To create an SSL gateway using the Configuration utility

1. In the navigation pane, click **Proxies**.
   The Proxies screen opens.
2. Click the **Add** button.
   The Add Proxy screen opens.
3. In the Add Proxy screen, configure the attributes you want to use with the proxy. For additional information about configuring a Proxy, click the **Help** button.

Configuration note

*For this example, create the following proxies on **Accelerator1**, **Accelerator2**, **Accelerator3**, and **Accelerator4**, respectively: **10.1.0.111:443**, **10.1.0.112:443**, **10.1.0.113:443**, and **10.1.0.114:443**.*

To create an SSL gateway from the command line

Use the following command syntax to create an SSL gateway:

```
b proxy <ip>:<service> \
target server <ip>:<service> \
clientssl enable \
clientssl key <clientssl_key> \
clientssl cert <clientssl_cert>
```

For example, to create the SSL gateways **accelerator1**, **accelerator2**, **accelerator3** and **acceslerator4**, you would use the following commands on these four e-Commerce Controllers, respectively. Note that the target for each gateway is the HTTP virtual server **192.168.200.30:80**. In this example, to complete the configuration for **accelerator1**, type the following command on **accelerator1**:

```
b proxy 10.1.0.111:443 \
target server 192.168.200.30:80 \
clientssl enable \
clientssl key my.server.net.key \
clientssl cert my.server.net.crt
```

In this example, to complete the configuration for **accelerator2**, type the following command on **accelerator2**:

```
b proxy 10.1.0.112:443 \
target server 192.168.200.30:80 \
clientssl enable \
clientssl key my.server.net.key \
clientssl cert my.server.net.crt
```

In this example, to complete the configuration for **accelerator3**, type the following command on **accelerator3**:

```
b proxy 10.1.0.113:443 \
target server 192.168.200.30:80 \
clientssl enable \
clientssl key my.server.net.key \
clientssl cert my.server.net.crt
```

In this example, to complete the configuration for **accelerator4**, type the following command on **accelerator4**:

```
b proxy 10.1.0.114:443 \
target server 192.168.200.30:80 \
clientssl enable \
clientssl key my.server.net.key \
clientssl cert my.server.net.crt
```

Enabling port 443

For security reasons, the ports on the SSL Accelerators do not accept traffic until you enable them. In this configuration, the SSL Accelerator accepts traffic on port **443** for SSL. For this configuration to work, you must enable port **443**. Use the following command to enable this port:

```
b service 443 tcp enable
```

Using SSL-to-server

As described so far, SSL Acceleration offloads SSL from the server to the BIG-IP. In some situations, security requirements demand that traffic on the internal VLAN (that is, behind the virtual server) be encrypted as well, or more exactly, re-encrypted. This server-side re-encryption requires that the servers handle the final SSL processing, but SSL acceleration is still obtained because the process is faster than allowing SSL client connections directly to the servers. (This is because session keys are re-used and because more efficient ciphers are used for the server-side SSL connections.) Figure 1.5 shows the SSL Accelerator configuration of Figure 1.1 with SSL-to-server added. Note that the only diagrammatic difference is that both client-side and server-side traffic are now labeled **SSL**, and the virtual server is now configured for service **443**.

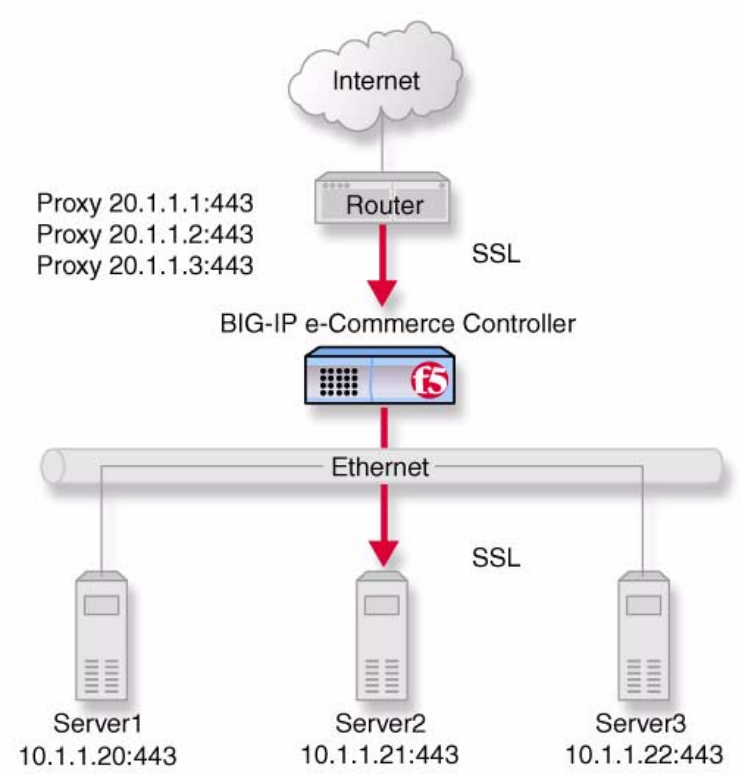

Figure 1.5 *An incoming SSL connection with SSL-to-server*

Configuring an SSL Accelerator with SSL-to-server

Since SSL-to-server is typically used together with standard, client-side SSL acceleration, configuring SSL-to-server involves the same tasks used in the preceding solutions (*Configuring the SSL Accelerator*, on page 1-2 and *Introducing the SSL Accelerator scalable configuration*, on page 1-9), with the following exceptions:

- The servers must be equipped and enabled for SSL processing.
- You must configure the server pool virtual server as HTTPS rather than HTTP and change the proxy targets accordingly.
- For the proxy or proxies, you must enable server-side SSL as well as the standard client-side SSL.

Optionally, you may configure a second certificate on the proxy to authenticate it to the servers as a trusted client.

Configuring a server pool and virtual server for HTTPS

To configure the server pool and virtual server for HTTPS for the non-scalable configuration, simply perform the steps in *Creating an SSL gateway*, on page 1-8 , but specify **20.1.1.10.443** as the target virtual server

and enable the **serverssl** attribute in addition to the **ssl** attribute. Entered from the command line, this would be accomplished as follows. In this example, to complete the configuration on the BIG-IP e-Commerce Controller, type the following command to create the proxy for **Server1**:

```
b proxy 20.1.1.1:443 \
target server 10.1.1.20:443 \
clientssl enable  \
clientssl key my.server.net.key \
clientssl cert my.server.net.crt \
serverssl enable
```

In this example, to complete the configuration on the BIG-IP e-Commerce Controller, type the following command to create the proxy for **Server2**:

```
b proxy 20.1.1.2:443 \
target server 20.1.1.21:443 \
clientssl enable \
clientssl key my.server.net.key \
clientssl cert my.server.net.crt
```

In this example, to complete the configuration on the BIG-IP e-Commerce Controller, type the following command to create the proxy for **Server3**:

```
b proxy 20.1.1.3:443 \
target server 10.1.1.22:443 \
clientssl enable \
clientssl key my.server.net.key \
clientssl cert my.server.net.crt
```

Optionally, you may specify a key file and a certificate file for the proxy as a client. This is done as follows:

```
b proxy 20.1.1.1:443 \
target server 10.1.1.20:443 \
clientssl enable \
clientssl key my.server.net.key \
clientssl cert my.server.net.crt \
serverssl enable \
serverssl key my.client.net.key \
serverssl cert my.client.net.key
```

Glossary

BIG-IP web server

The BIG-IP web server runs on a BIG-IP and hosts the Configuration utility.

bigpipe

The **bigpipe** utility provides command line access to the BIG-IP.

BIG/stat

BIG/stat is a statistical monitoring utility that ships on the BIG-IP. This utility provides a snap-shot of statistical information.

BIG/top

BIG/top is a statistical monitoring utility that ships on the BIG-IP. This utility provides real-time statistical information.

big3d

The **big3d** agent is a monitoring utility that collects metrics information about paths between a BIG-IP and a specific local DNS server. The **big3d** agent runs on BIG-IP units and it forwards metrics information to 3-DNS systems.

BIND (Berkeley Internet Name Domain)

BIND is the most common implementation of DNS, which provides a system for matching domain names to IP addresses.

chain

A chain is a series of filtering criteria used to restrict access to an IP address. The order of the criteria in the chain determines how the filter is applied, from the general criteria first, to the more detailed criteria at the end of the chain.

default VLANs

The BIG-IP is configured with two default VLANs, one for each interface. One default VLAN is named **internal** and one is named **external**. See also *VLAN*.

dynamic site content

Dynamic site content is site content that is automatically generated each time a user accesses the site. Examples are current stock quotes or weather satellite images.

external VLAN

The external VLAN is a default VLAN on the BIG-IP. In a basic configuration, this VLAN has the administration ports locked down. In a normal configuration, this is typically a VLAN on which external clients request connections to internal servers.

floating self IP address

A floating self IP address is an additional self IP address for a VLAN that serves as a shared address by both units of a BIG-IP redundant system.

health check

A health check is a BIG-IP feature that determines whether a node is **up** or **down**. Health checks are implemented through health monitors. See also *health monitor*.

health monitor

A health monitor checks a node to see if it is **up** and functioning for a given service. If the node fails the check, it is marked **down**.

HTTP redirect

An HTTP redirect sends an HTTP 302 Object Found message to clients. You can configure a pool with an HTTP redirect to send clients to another node or virtual server if the members of the pool are marked **down**.

ICMP (Internet Control Message Protocol)

ICMP is an Internet communications protocol used to determine information about routes to destination addresses, such as virtual servers managed by BIG-IP and 3-DNS systems.

interface

The physical port on a BIG-IP is called an interface. See also *link*.

internal VLAN

The internal VLAN is a default VLAN on the BIG-IP. In a basic configuration, this VLAN has the administration ports open. In a normal configuration, this is a network interface that handles connections from internal servers.

iQuery

A UDP based protocol used to exchange information between BIG-IP units and 3-DNS systems. The iQuery protocol is officially registered for port 4353.

last hop

A last hop is the final hop a connection took to get to the BIG-IP. You can allow the BIG-IP to determine the last hop automatically to send packets back to the device from which they originated.

link

A link is a physical interface on the BIG-IP connected to another physical interface in a network.

link aggregation

The link aggregation feature allows you to combine a number of links together to act as one interface.

loopback adapter

A loopback adapter is a software interface that is not associated with an actual network card. The nPath routing configuration requires you to configure loopback adapters on servers.

MAC (Media Access Control)

MAC is a protocol that defines the way workstations gain access to transmission media, and is most widely used in reference to LANs. For IEEE LANs, the MAC layer is the lower sublayer of the data link layer protocol.

MAC address

A MAC address is used to represent hardware devices on an Ethernet network.

monitor

The BIG-IP uses monitors to determine whether nodes are **up** or **down**. There are several different types of monitors and they use various methods to determine the status of a server or service.

monitor instance

You create a monitor instance when a health monitor is associated with a node, node address, or port. It is the monitor instance that actually performs the health check, not the monitor.

monitor template

A monitor template is a system-supplied health monitor that is used primarily as a template to create user-defined monitors, but in some cases can be used as is. The BIG-IP includes a number of monitor templates, each

specific to a service type, for example, HTTP and FTP. The template has a template type that corresponds to the service type and is usually the name of the template.

named

Named is the name server utility, which manages domain name server software.

NAT (Network Address Translation)

A NAT is an alias IP address that identifies a specific node managed by the BIG-IP to the external network.

node

A node is a specific combination of an IP address and port (service) number associated with a server in the array that is managed by the BIG-IP.

node address

A node address is the IP address associated with one or more nodes. This IP address can be the real IP address of a network server, or it can be an alias IP address on a network server.

node alias

A node alias is a node address that the BIG-IP uses to verify the status of multiple nodes. When the BIG-IP uses a node alias to check node status, it pings the node alias. If the BIG-IP receives a response to the ping, it marks all nodes associated with the node alias as **up**. If the controller does not receive a response to the ping, the it marks all nodes associated with the node alias as **down**.

node port

A node port is the port number or service name that is hosted by a specific node.

node status

Node status indicates whether a node is **up** and available to receive connections, or **down** and unavailable. The BIG-IP uses the node ping and health check features to determine node status.

persistence

A series of related connections received from the same client, having the same session ID. When persistence is turned on, a controller sends all connections having the same session ID to the same node instead of load balancing the connections.

port

A port is represented by a number that is associated with a specific service supported by a host. Refer to the Services and Port Index for a list of port numbers and corresponding services.

rate class

You create a rate filter from the Configuration utility or command line utility. When you assign a rate class to a rate filter, a rate class determines the volume of traffic allowed through a rate filter. See also *rate filter*.

rate filter

Rate filters consist of a basic filter with a rate class. Rate filters are a type of extended IP filter. They use the same IP filter method, but they apply a rate class, which determines the volume of network traffic allowed through the filter. See also *rate class*.

remote administrative IP address

The remote administrative IP address is an IP address from which a controller allows shell connections, such as Telnet or SSH.

RFC 1918 addresses

An RFC 1918 address is an IP address that is within the range of non-routable addresses described in the IETF RFC 1918.

self IP address

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

service

Service refers to services such as TCP, UDP, HTTP, and FTP.

Setup utility

The Setup utility walks you through the initial system configuration process. You can run the Setup utility from either the command line or the Configuration utility start page.

SNAT (Secure Network Address Translation)

A SNAT is a feature you can configure on the BIG-IP. A SNAT defines a routable alias IP address that one or more nodes can use as a source IP address when making connections to hosts on the external network.

SNAT automap

This feature allows the BIG-IP to perform a SNAT automatically on any connection that is coming from the controller's internal VLAN. It is easier to use than traditional SNATs and solves certain problem associated to traditional SNATs.

SNMP (Simple Network Management Protocol)

SNMP is the Internet standard protocol, defined in STD 15, RFC 1157, developed to manage nodes on an IP network.

SSL gateway

An SSL gateway is a gateway for decrypting HTTPS requests to an HTTP server and encrypting the reply.

state mirroring

State mirroring is a feature on the BIG-IP that preserves connection and persistence information in a BIG-IP redundant system.

static load balancing modes

Static load balancing modes base connection distribution on a pre-defined list of criteria; it does not take current server performance or current connection load into account.

sticky mask

A sticky mask is a special IP mask that you can configure on the BIG-IP. This mask optimizes sticky persistence entries by grouping more of them together.

tagged VLAN

You can define any interface as a member of a tagged VLAN. You can create a list of VLAN tags or names for each tagged interface.

transparent node

A transparent node appears as a router to other network devices, including the BIG-IP.

trunk

A trunk is a combination of two or more interfaces and cables configured as one link. See also *link aggregation.*

virtual address

A virtual address is an IP address associated with one or more virtual servers managed by the BIG-IP.

virtual port

A virtual port is the port number or service name associated with one or more virtual servers managed by the BIG-IP. A virtual port number should be the same TCP or UDP port number to which client programs expect to connect.

virtual server

Virtual servers are a specific combination of virtual address and virtual port, associated with a content site that is managed by a BIG-IP or other type of host server.

VLAN

VLAN stands for virtual local area network. A VLAN is a logical grouping of network devices. You can use a VLAN to logically group devices that are on different network segments.

VLAN name

A VLAN name is the symbolic name used to identify a VLAN. For example, you might configure a VLAN named marketing, or a VLAN named development. See also *VLAN*.

Index

G

H

I

K

L

M

N

O

P

U

V

W

X