



BIG-IP[®] Network and System Management Guide

version 9.0 and later

Product Version

This manual applies to version 9.0 and later of the BIG-IP® product family.

Legal Notices

Copyright

Copyright © 1996-2005, F5 Networks, Inc. All rights reserved.

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 by applicable iControl user licenses. F5 reserves the right to change specifications at any time without notice.

Trademarks

F5, F5 Networks, the F5 logo, BIG-IP, 3-DNS, iControl, FireGuard, Internet Control Architecture, IP Application Switch, iRules, OneConnect, Packet Velocity, SYN Check, Control Your World, ZoneRunner, uRoam, FirePass, and TrafficShield are registered trademarks or trademarks of F5 Networks, Inc. in the U.S. and certain other countries. All other trademarks mentioned in this document are the property of their respective owners. F5 Networks' trademarks may not be used in connection with any product or service except as permitted in writing by F5.

Patents

This product protected by U.S. Patents 6,374,300; 6,473,802. Other patents pending.

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 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

1

Introducing BIG-IP Network and System Management

Introducing the BIG-IP system	1-1
Overview of BIG-IP network management features	1-1
Overview of BIG-IP system management features	1-2
Getting started	1-5
Choosing a configuration tool	1-5
Finding additional help and technical support resources	1-6
Understanding stylistic conventions	1-7

2

Operating the LCD Panel

Introducing the LCD panel	2-1
Using the LCD panel	2-2
Pausing on a screen	2-2
Using LCD menus	2-2
Powering up the unit	2-2
Halting the unit	2-2
Powering down the unit	2-3
Rebooting the unit	2-3
Clearing alerts	2-3
Navigating through the LCD menus	2-4

3

Lights-Out Management

Introducing lights-out management	3-1
Accessing the command menu	3-3
Options for accessing the command menu	3-3
Setting up remote lights-out SSH access	3-4
Using out-of-band management commands	3-4
Using the command menu	3-5

4

Configuring the BIG-IP Platform and General Properties

Introducing the BIG-IP platform and general properties	4-1
Configuring platform properties	4-2
Configuring platform properties and user administration settings	4-2
Managing a web server certificate	4-5
Configuring general properties	4-9
Configuring device-related properties	4-9
Configuring local-traffic properties	4-13

5

Working with Interfaces

Introducing BIG-IP system interfaces	5-1
Configuring interfaces	5-2
Understanding interface naming conventions	5-2
Viewing interface information and media properties	5-2
Configuring interface properties	5-3
Configuring interface mirroring	5-6
Displaying interface statistics	5-7
Related configuration tasks	5-9

6

Configuring Self IP Addresses

Introducing self IP addresses	6-1
Types of self IP addresses	6-1
Self IP addresses and MAC addresses	6-2
Using self IP addresses for SNATs	6-2
Creating and managing self IP addresses	6-3
Creating a self IP address	6-3
Managing self IP addresses	6-6

7

Configuring VLANs and VLAN Groups

Introducing virtual LANs	7-1
Understanding VLANs on a BIG-IP system	7-1
Understanding the default VLAN configuration	7-2
Creating and managing VLANs	7-3
Creating a VLAN	7-3
Managing a VLAN	7-10
Creating and managing VLAN groups	7-13
Creating a VLAN group	7-14
Managing a VLAN group	7-18
Assigning self IPs to VLANs and VLAN groups	7-19

8

Working with Trunks

Introducing trunks	8-1
How do trunks work?	8-1
Overview of LACP	8-2
Creating a trunk	8-3
Specifying a trunk name	8-4
Specifying interfaces for a trunk	8-4
Enabling LACP	8-5
Specifying the LACP mode	8-5
Specifying the LACP timeout	8-6
Specifying a link selection policy	8-6
Managing trunks	8-9
Viewing a list of trunks	8-9
Viewing or modifying trunk properties	8-9
Adding a trunk to a VLAN	8-10
Deleting a trunk	8-10
Managing interfaces for a trunk	8-11

9

Configuring Spanning Tree Protocols

Introducing spanning tree protocols	9-1
Spanning tree protocol types	9-1
Using spanning tree with legacy bridges	9-3
Configuration overview	9-4
Configuring global spanning tree properties	9-5
Specifying the spanning tree mode	9-6
Configuring global timers	9-7
Specifying the Transmit Hold Count option	9-8
Configuring MSTP-specific global properties	9-8

Managing spanning tree instances	9-10
Viewing a list of spanning tree instances	9-10
Creating a spanning tree instance (MSTP-only)	9-11
Viewing and modifying a spanning tree instance	9-14
Deleting a spanning tree instance or its members (MSTP-only)	9-15
Configuring interfaces for spanning tree	9-16
Configuring spanning tree settings on an interface	9-16
Managing interfaces for a specific instance	9-19

10

Managing User Accounts

Introducing user account management	10-1
Understanding user account types	10-1
Understanding user roles	10-2
Managing local user accounts	10-4
Configuring the admin account	10-4
Creating user accounts	10-4
Viewing user accounts	10-5
Modifying user accounts	10-6
Deleting user accounts	10-7
Managing remote user accounts	10-8
Specifying a remote user-account database	10-8
Assigning user roles	10-9
Deleting an explicit user-role designation	10-11

11

Configuring SNMP

Introducing SNMP administration	11-1
Reviewing an industry-standard SNMP implementation	11-1
Reviewing the BIG-IP system SNMP implementation	11-1
Summarizing SNMP configuration on the BIG-IP system	11-2
Configuring the SNMP agent	11-3
Configuring BIG-IP system information	11-3
Configuring client access	11-4
Controlling access to SNMP data	11-5
Configuring traps	11-7
Working with SNMP MIB files	11-9
Downloading SNMP MIB files	11-9
Understanding the enterprise MIB files	11-10
Collecting performance data	11-13
Collecting data on memory use	11-14
Collecting data on active connections	11-14
Collecting data on new connections	11-15
Collecting data on throughput	11-16
Collecting data on HTTP requests	11-16
Collecting data on CPU use	11-17

12

Setting up a Redundant System

Introducing redundant systems	12-1
Summary of redundant system features	12-1
Understanding failover and fallback	12-2
Understanding self IP addresses for redundant systems	12-6

Understanding fail-safe	12-8
Before you begin	12-8
Configuring units of a redundant pair	12-11
Specifying primary and secondary failover addresses	12-12
Configuring the redundancy mode	12-13
Specifying a redundancy state preference	12-14
Configuring the failover type	12-15
Specifying the link down-time on failover	12-15
Specifying the default route on back-end servers	12-16
Special considerations for active-active systems	12-16
The next step	12-17
Synchronizing configuration data	12-18
Performing configuration synchronization	12-19
Enabling the global display of synchronization status	12-20
Continuing with active-active system configuration	12-21
Associating BIG-IP system objects with unit IDs	12-21
Synchronizing the configuration	12-23
Configuring fail-safe	12-23
Configuring system fail-safe	12-23
Configuring gateway fail-safe	12-25
Configuring VLAN fail-safe	12-26
Mirroring connection information	12-28
Setting a shared MAC masquerade address	12-29
Viewing interfaces and MAC addresses	12-29
Designating a shared MAC masquerade address	12-30
Maintaining a redundant system	12-32
Viewing redundancy states and synchronization status	12-32
Changing the redundancy state	12-35
Determining the unit ID	12-35
Customizing redundant-system behavior	12-36
Controlling failback	12-36
Converting an active-active system to an active/standby system	12-38

A

Troubleshooting SNMP Traps

Understanding F5-specific traps	A-1
General traps	A-1
Hardware-related traps	A-2
License-related traps	A-2
TMOS-related traps	A-3
Authentication-related traps	A-4
DoS-related traps	A-4
Network-related traps	A-4
Logging-related traps	A-5

B

Configuring bigdb Database Keys

Introducing the bigdb database	B-1
Summarizing bigdb keys for redundant system administration	B-2
Using failover keys	B-2
Using connection mirroring keys	B-4
Using configuration synchronization keys	B-5
Using system fail-safe keys	B-6
Summarizing bigdb keys for user account administration	B-7

Glossary

Index

Table of Contents



I

Introducing BIG-IP Network and System Management

- Introducing the BIG-IP system
- Getting started

Introducing the BIG-IP system

BIG-IP system is a port-based, multilayer switch. Using a BIG-IP system means that you can reduce the need for routers and IP routing by managing traffic at the data-link layer (layer 2). This in turn reduces equipment costs and boosts overall network performance. At the same time, the BIG-IP system's multilayer capabilities allow the system to process traffic at other OSI layers when it is advantageous to do so. The BIG-IP can perform IP routing at layer 3 when needed, as well as manage TCP and application traffic at layers 4 and 7.

Overview of BIG-IP network management features

The BIG-IP system consists of several fundamental network components that you can configure in the way that best utilizes BIG-IP system capabilities.

Interfaces, spanning tree protocols, and trunks

A BIG-IP system has either four, eight, or sixteen interfaces for switching or routing traffic from various hosts or other devices on the network.

Interfaces are the hardware ports that the BIG-IP system uses to send and receive traffic. For more information, see Chapter 5, *Working with Interfaces*.

When you connect multiple switches to the BIG-IP system in parallel, you can configure your hosts to make use of spanning tree protocols. **Spanning tree protocols** provide path redundancy while preventing unwanted loops in the network. You can view spanning tree instances, configure global spanning tree options, and configure spanning tree settings for each interface. For optimal performance, you can use spanning tree protocols in conjunction with the trunks feature. For more information, see Chapter 9, *Configuring Spanning Tree Protocols*.

Trunks are a feature you can use to aggregate your links. When you create **trunks**, you group interfaces together to function as one larger interface and to provide redundancy if one interface in the trunk becomes unavailable. When that occurs, traffic can be processed on another interface in the trunk. For more information, see Chapter 8, *Working with Trunks*.

VLANs and self IP addresses

A virtual local area network, or **VLAN**, is a logical collection of hosts on the network. Each VLAN has one or more BIG-IP system interfaces associated with it. VLANs have these primary advantages:

- ◆ **VLANs define boundaries for a broadcast domains.**

Traditionally, network administrators have deployed routers within the same IP network, to define smaller broadcast boundaries. A better solution is to use VLANs. When a host in a VLAN sends a broadcast

message to find the MAC address of a destination host, the message is sent to only those hosts in the VLAN. Using VLANs to control the boundaries of broadcast domains prevents messages from flooding the network, thus enhancing network performance.

◆ **VLANs ease system and network maintenance**

Normally, the way to enable hosts to share network resources, such as storage devices and printers, has been to group hosts into the same physical location. Continually moving and re-cabling hosts to other locations on the network, as well as manually updating routing tables, can be a costly and time-consuming task for a system or network administrator. Using VLANs, you can avoid these problems. All hosts that you group within a VLAN can share network resources, regardless of their physical location on the network.

To enhance performance and flexibility, the BIG-IP system comes with two existing virtual local area networks (VLANs), one for your external network, and one for your internal network. Each of these VLANs has a BIG-IP system interface already assigned to it. You can use these two VLANs as is, you can assign additional interfaces to these VLANs, or you can create more VLANs and assign interfaces to them, too. For more information, see Chapter 7, *Configuring VLANs and VLAN Groups*.

Each VLAN you create has its own *self IP address*. The BIG-IP system uses this address as the source IP address when sending requests to hosts in a VLAN, and hosts in a VLAN use this IP address as the destination IP address when sending responses to the BIG-IP system.

When you first ran the Setup utility, you assigned a self IP address to the internal VLAN, and another self IP address to the external VLAN. As you create other VLANs, you assign self IP addresses to them, too. Also, units of a redundant system can share a self IP address, to ensure that the BIG-IP system can process server responses successfully when failover has occurred. For more information, see Chapter 6, *Configuring Self IP Addresses*.

Overview of BIG-IP system management features

This guide addresses some of the system management options that are common to all BIG-IP systems. These options include creating and maintaining administrative user accounts, configuring System Network Management Protocol (SNMP), and configuring and maintaining redundant systems.

You partially configure some of these options by running the Setup utility on the BIG-IP system. Once you have run the Setup utility, you can use the Configuration utility to complete the configuration of these options and to manage the BIG-IP system on an ongoing basis.

Liquid crystal display and lights-out management

Using the *liquid crystal display (LCD)*, you can control the BIG-IP unit without attaching a serial or network cable. With the *lights out management* feature, you can remotely manage certain aspects of the operation of the hardware unit and the BIG-IP traffic management operating system in the event that the traffic management software becomes incapacitated. For more information, see Chapter 2, *Operating the LCD Panel*, and Chapter 3, *Lights-Out Management*.

User accounts

You can create or manage user accounts for BIG-IP system administrators. These accounts can reside either locally on the BIG-IP system, or remotely on a separate authentication server such as a Lightweight Directory Access Protocol (LDAP) server. You can also manage the three special user accounts **root**, **admin**, and **support**.

For each new user account that you create, you can assign a user role that defines the type and level of access granted to that user. The available user roles are: **Administrator**, **Operator**, **Guest**, and **No Access**.

The types of remote authentication servers that you can use to store user accounts for BIG-IP system administrators are: Active Directory™ servers, Lightweight Directory Access Protocol (LDAP) servers, and Remote Authentication Dial-in User Service (RADIUS) servers. For more information, see Chapter 10, *Managing User Accounts*.

System Network Management Protocol (SNMP)

System Network Management Protocol (SNMP) is an industry-standard protocol that allows you to manage the BIG-IP system remotely, along with other devices on the network. The BIG-IP system provides the SNMP agent and the MIB files that you need to manage the system remotely using SNMP. For more information, see Chapter 11, *Configuring SNMP*.

Redundant systems

To ensure high-availability of the BIG-IP system, you can set up a redundant-system configuration. Then, if one BIG-IP system becomes unavailable, another BIG-IP system can immediately take over to process the traffic.

When you first run the Setup utility on a BIG-IP system, you specify whether the system is a unit of a redundant pair. When you configure two BIG-IP systems to function as units of a redundant system, a process known as failover occurs when one of those units becomes unavailable for any reason. *Failover* ensures that the BIG-IP system can still process traffic when a unit is unavailable.

Every redundant system has a mode that you specify, either active/standby or active-active. If you choose active/standby mode and failover occurs later, then by default the standby unit becomes active, and remains active, until failover occurs again. If you choose active-active mode, the surviving unit begins processing connections targeted for the failed unit, while continuing to process its own connections. In this way, users experience no interruption in service in the event of system unavailability. For more information, see Chapter 12, *Setting up a Redundant System*.

Getting started

Before you use this guide, we recommend that you run the Setup utility on the BIG-IP system to configure basic network elements such as static and floating self IP addresses, interfaces, and VLANs, to name a few. After running the Setup utility, you can further customize your system by using the Configuration utility to configure local traffic management objects such as virtual servers, load balancing pools, and profiles. Finally, you can return to this guide to configure and maintain the features introduced in *Introducing the BIG-IP system*, on page 1-1.

Before you use this guide, we recommend that you do the following:

- Choose a configuration tool.
- Familiarize yourself with additional resources such as product guides and online help.
- Review the stylistic conventions that appear in this chapter.

Choosing a configuration tool

The BIG-IP system offers a browser-based utility for managing the BIG-IP system, and, as an alternative, various command line utilities. Note that all procedures in this guide describe how to manage the system using the browser-based utility.

The Configuration utility

The Configuration utility is a browser-based application that you use to configure and monitor the load balancing setup on the BIG-IP system. Once you complete the instructions for the Setup utility described in this guide, you can use the Configuration utility to perform additional 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 MIBs or the SSH client. The Configuration utility requires Netscape Navigator version 4.7, or Microsoft Internet Explorer version 5.0 or 5.5.

Command-line utilities

In addition to using the Configuration utility, you can also manage the BIG-IP system using command-line utilities such as the **bigpipe** utility™. To monitor the BIG-IP system, 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 system console, or you can run commands using a remote shell, such as the SSH client or a Telnet client. For more information on command-line utilities, see the online man pages.

Finding additional help and technical support resources

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

◆ **BIG-IP system product guides**

The Ask F5 site offers PDF versions of these guides:

- ***Platform Guide: 1500, 3400, 6400, and 6800***
This guide includes information about the BIG-IP unit. It also contains important environmental warnings.
- ***Installation, Licensing, and Upgrades for BIG-IP Systems***
This guide includes information you need to initially install or upgrade the BIG-IP system.
- ***Configuration Guide for Local Traffic Management***
This guide provides detailed step-by-step procedures for configuring the local traffic management system, using the browser-based Configuration utility.
- ***System Management Guide*** (this guide)
This guide provides concepts and detailed step-by-step procedures for maintaining certain parts of the BIG-IP system such as user accounts and SNMP, using the browser-based Configuration utility.
- ***New and Updated Commands***
This guide provides information about various command-line utilities, such as the **bigpipe** utility, that you can use to configure and manage a BIG-IP system.

◆ **Online help**

In addition to the product guides, the following types of help are available online:

- **Online solutions**
Through our Ask F5 Web site, you can view examples of common load balancing solutions. Before you begin installing the hardware, we recommend that you browse this site to find the load balancing solution that works best for you.
- **Online help screens**
The browser-based Configuration utility includes online help for each screen. Simply click the **Help** button.
- **man pages**
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 system displays the syntax and usage associated with the command.

Understanding stylistic conventions

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

Using the solution examples

All examples in this document use only private class IP addresses. When you set up the solutions we describe, you must use valid 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 ***floating IP address*** is an IP address assigned to a VLAN and shared between two computer systems.

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 self <ip_address> show** command, you can specify a specific self IP address to show by specifying an IP address for the **<ip_address>** 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 SNMP traps in the ***Network and System Management Guide***, Appendix A, *Troubleshooting SNMP Traps*.

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 self <ip_address> show
```

or

```
b self <ip_Address> show
```

Table 1.1 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 1.1 *Command line syntax conventions*



2

Operating the LCD Panel

- Introducing the LCD panel
- Using the LCD panel
- Navigating through the LCD menus

Introducing the LCD panel

The liquid crystal display, or LCD panel, provides the ability to control the unit without attaching a serial or network cable. The following menus are available on the LCD panel.

- ◆ **Information menu**

Use the Information menu to find information about using the LCD and its functionality.

- ◆ **System menu**

Use the System menu to reboot, netboot, or halt the unit. This menu also has options for setting the properties of the management interface (MGMT) and the serial port.

- ◆ **Screen menu**

Use the Screen menu to set up the informational screens you would like the LCD to cycle through. The information screens include system status, statistics, and system alerts.

- ◆ **Options menu**

Use the Options menu to configure the properties of the LCD panel.

This chapter describes how to use the LCD panel and its menus. It does not describe each function available in each menu.

Figure 2.1 shows an example of the LCD panel and control buttons.



Figure 2.1 An example of the LCD panel and control buttons

Using the LCD panel

You can configure the LCD panel to meet your needs. The following section describes how to perform a number of tasks with the LCD panel:

- Pause on a screen
- Use the LCD menus
- Power up the unit
- Halt the unit
- Power down the unit
- Reboot the unit

Pausing on a screen

Normally, the screens cycle on the LCD at a constant rate. However, push the Check button to toggle the LCD between Hold and Rotate modes. In Hold mode, a single screen is displayed. The Rotate mode changes the screen displayed on the LCD every 4 seconds.

Using LCD menus

Pressing the **X** button puts the LCD panel in Menu mode. The buttons Left Arrow, Right Arrow, Up Arrow, and Down Arrow are only functional when the LCD is in Menu mode.

Powering up the unit

When you want to power on a unit that is shut down, press the Check button to turn the power on.

Halting the unit

We recommend you halt the unit before you power it down or reboot it using the LCD menu options.

To halt the unit

1. Press the **X** button, then use the arrow keys to navigate to the System menu.
2. Press Check. Navigate to the Halt menu.
3. Press the Check button. Press the Check button again at the confirmation screen.
4. Wait 30 seconds before powering the machine off or rebooting it.

Powering down the unit

Hold the **X** button for 4 seconds to power down the unit. We recommend that you halt the system before you power down the system in this manner.

Rebooting the unit

Hold the **Check** button for 4 seconds to reboot the unit. You should only use this option after you halt the unit.

Clearing alerts

Press the **Check** button to clear any alerts on the LCD screen. You must clear any alerts on the screen before you can use the LCD.

Navigating through the LCD menus

To use the LCD menus, you must first put the LCD in menu mode. To put the LCD in menu mode, press the **X** button.

After you put the LCD in menu mode, use the Left Arrow, Right Arrow, Up Arrow, and Down Arrow buttons to select menu options. There are four menu options:

- Information
- System
- Screens
- Options

The following tables describe each LCD menu option.

Information menu

You can use the Information menu to access help pages about using the LCD panel functionality. You can also find more information on what different LED activity means, and the failover state of the unit in a redundant pair. The following table, Table 2.1, shows the options available on the Information menu.

Option	Description
How to use the LCD	Displays a vertical scrolling text description on how to use the LCD panel.
Front Panel LEDs	Displays a vertical scrolling text description of what the front panel LEDs mean.
Port Indicators	Displays a vertical scrolling text description of what the lights above the ports mean.
Console and Failover serial port information	Displays a vertical scrolling text description of the console and failover serial ports.

Table 2.1 *The Information menu*

System menu

The System menu provides various options for rebooting, halting, or netbooting the hardware. This menu also provides options for configuring the network on the management interface. The following table, Table 2.2, lists the options available in the System menu.

Option	Description
Reboot	Select this option to reboot the unit.
Halt	Select this option to halt the unit.
Netboot	Select this option if you are installing software from a PXE server.
IP address	Type the management interface IP address. You can only use an IPv4 address.
Netmask	Set the netmask for the management interface IP address.
Default route	Type in the default route for the management interface. This route is necessary if you plan to manage the unit from a different subnetwork.
Commit	Select this option to commit your changes.
Serial port	Use this option to change the baud rate of the serial port. The following options are available: 9600 19200 38400 115200

Table 2.2 The System menu

Screens menu

You can use the Screens menu options to view various statistics and information about the system. The following table, Table 2.3, lists all the general information screens. You can use the LCD button to place a check mark next to the name of the screens you would like to appear when the screens cycle.

Option	Description
Version screen	Displays the product version information.
Information screen	Displays the information screen menu.
Date and Time screen	Displays the date and time.
MAC addresses screen	Displays the MAC addresses on the unit.

Table 2.3 The general screen information menu

Option	Description
System information screen	Displays system information.
CPU usage	Displays the CPU usage percentage.
Memory usage	Displays the memory usage.
Auth requests	Displays the number of authentication requests being processed.
Statistics	Displays simple statistics, such as bytes and packets in and out of the system.
Alert screen	Displays system alerts.

Table 2.3 The general screen information menu

Options menu

You can use the Options menu to adjust the display properties of the LCD panel. The following table, Table 2.4, lists the options available on the Options menu.

Option	Description
Contrast	Use the Left and Right arrow keys on the LCD to set the contrast of the LCD.
On Brightness	This setting provides the ability to adjust the LCD backlight brightness.
Off Brightness	This setting controls the brightness of the LCD panel when the backlight is off. Use the Left and Right arrow keys to set the brightness of the LCD panel.

Table 2.4 The Options menu



3

Lights-Out Management

- Introducing lights-out management
- Accessing the command menu

Introducing lights-out management

A lights-out management system is available with the latest F5 Networks platforms. The *lights-out* management system provides the ability to remotely manage certain aspects of the operation of the hardware unit and the BIG-IP traffic management operating system in the event the traffic management software becomes incapacitated.

The lights-out management system consists of the following elements.

- ◆ **Switch card control processor (SCCP)**
The hardware that provides the hardware control over the whole unit.
- ◆ **Host console shell (hostconsh)**
The shell that provides access to the command menu.
- ◆ **Command menu**
The menu that contains the options for lights-out management.
- ◆ **Traffic management operating system**
The software that you configure to manage the traffic for your site.
- ◆ **Out-of-band management commands**
The commands that provide the ability to control various aspects of the system with a series of keystrokes.

The command menu operates independently of the traffic management operating system through the management port, the serial port console, and remotely through the traffic management ports.

- You can use the command menu to reset the unit, even if the BIG-IP traffic management system has locked up.
- You can remotely set a unit to netboot for a software re-install from an ISO image.
- You can get console access to the BIG-IP traffic management system itself, so you can configure the traffic management system from the command line interface.

The lights-out management system and the BIG-IP traffic management system function independently within the hardware unit. Figure 3.1 shows the relationship between the lights-out management system and the traffic management system.

The lights-out management system is accessible through the management interface (number 1 in Figure 3.1) and the console port (number 2 in Figure 3.1). This functionality is independent of the traffic management system (number 3 in Figure 3.1).

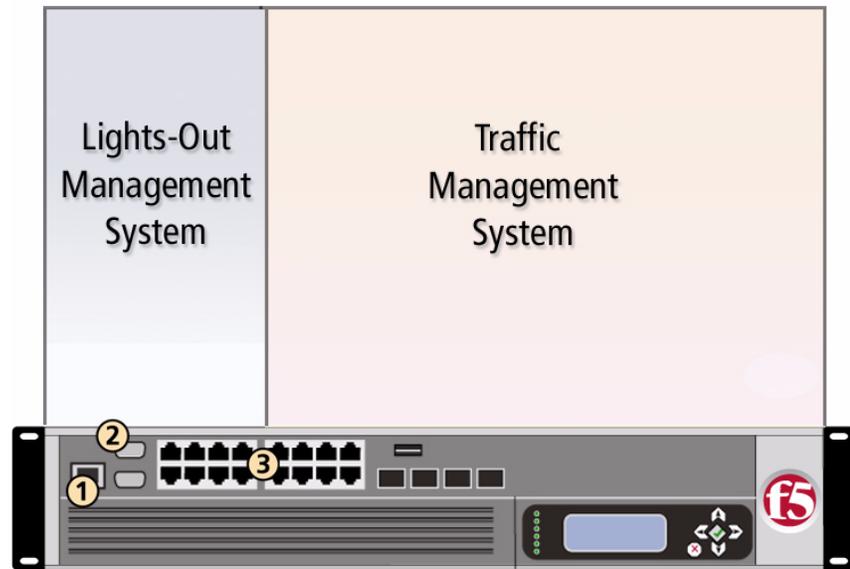


Figure 3.1 The lights-out management system and the traffic management system.

Accessing the command menu

You can access the command menu through the host console shell (**hostconsh**) using the front panel serial console, or remotely through SSH. The following section describes how to access the command menu both through the serial console and with an SSH client to the management interface.

Options for accessing the command menu

There are two methods you can use to access the command menu. You can access the menu from the serial console, or from an SSH client to the management interface.

To access the command menu from the serial console

1. From the serial console connected to the CONSOLE port, type the following key sequence.

Esc (

2. The command menu opens.
For details about each option on the command menu, see *Using the command menu*, on page 3-5.

To access the command menu using SSH

Before you can access the command menu using SSH, you must also have an IP address configured for remote lights-out management. For more information, see *Setting up remote lights-out SSH access*, on page 3-4.

1. Open the SSH client on a management workstation connected to the MGMT port on the BIG-IP system.
2. Type the following command, where **<IP addr>** is the IP address you configured for the lights-out system.

```
ssh console@<ip addr>
```

3. The host console shell opens.
4. To open the command menu, type the following key sequence.

Esc (

For details about each option on the command menu, see *Using the command menu*, on page 3-5.

Setting up remote lights-out SSH access

You can use the command menu to set up remote SSH access to the BIG-IP system. To set up remote access, run the SCCP network configuration utility to configure an IP address, netmask, and gateway for the lights-out system. You can only connect remotely with the SSH client through the management network connected to the management port (MGMT).

To configure remote SSH lights-out access

1. Log into the BIG-IP system through the serial console.
2. Type the following key sequence.
Esc (
3. After the command menu opens, type **N**.
This starts the network configuration utility for the SCCP.
4. Add an IP address, netmask, and gateway on the management network.

Using out-of-band management commands

The host console shell implements a subset of the Microsoft standard out-of-band management protocol. These commands provide the ability to use a series of key strokes to manage the host processor. Table 3.1 lists the key-stroke options that are available.

Keystroke combination	Result
Esc + R + Esc + r + Esc + R	You can use this sequence during pass-through mode to reboot the platform. We do not recommend using this method to reboot the platform.
Esc (Brings up the command menu.

Table 3.1 Out-of-band management key combinations

Using the command menu

The command menu provides the lights-out management options for the system (Figure 3.1).

```

1 --- Connect to Host subsystem console
2 --- Select Host subsystem boot mode: boot from local drive
3 --- Select Host subsystem boot mode: netboot from SCCP
4 --- Select Host subsystem boot mode: netboot from external server
5 --- Reboot Host subsystem (sends reboot command)
6 --- Halt Host subsystem (sends halt command)
7 --- Reset Host subsystem (issues hardware reset--USE WITH CARE!)
8 --- Reboot SCCP subsystem (issues hardware reset--USE WITH CARE!)
9 --- Halt SCCP subsystem (issues hardware shutdown--USE WITH CARE!)
B --- SCCP baud rate configurator
L --- SCCP login
N --- SCCP network configurator

```

Figure 3.2 A console view of the host processor console command menu

Each of these options is described in Table 3.2. Note that some of these commands are not intended for use by end users. Table 3.2 also specifies which commands are not recommended for use by users.

Option	Description
1	Exits the command menu and returns to terminal emulation mode.
2	Configures the BIG-IP traffic management system to boot from the local hard drive or CompactFlash card.
3	Configures the host subsystem to netboot from the host subsystem processor. This option is only for factory testing.
4	Configures the SCCP to netboot the host processor from an external server attached to the management network interface. This option provides the ability to start the PXE installation process remotely.
5	Reboots the host subsystem. In this case, the BIG-IP traffic management operating system (TMOS) is rebooted.
6	Halts the host subsystem. In this case, the BIG-IP traffic management operating system (TMOS) is halted.
7	Resets the host subsystem. In this case, the system is reset with a hardware reset.
8	Reboots the switch card control processor (SCCP). This resets the entire unit.
9	Halts the switch card control processor (SCCP). This shuts down the entire unit.

Table 3.2 Command menu options

Option	Description
B	Runs the switch card control processor (SCCP) baud rate configuration utility. This utility provides the ability to configure the SCCP serial speed and parameters. This option is only available through the front panel serial console.
L	Presents a login prompt for the switch card control processor (SCCP) subsystem. This subsystem cannot be configured by end users. This option is only available through the front panel serial console.
N	Runs the switch card control processor (SCCP) network configuration utility. This utility provides the ability to reconfigure the IP address, netmask, and default gateway used by the SCCP. If you change these settings, your session is disconnected. This option is only available through the front panel serial console.

Table 3.2 Command menu options

◆ Important

We do not recommend using the reset option, option 7, under normal circumstances. It does not allow for graceful shutdown of the BIG-IP system.



4

Configuring the BIG-IP Platform and General Properties

- Introducing the BIG-IP platform and general properties
- Configuring platform properties
- Configuring general properties

Introducing the BIG-IP platform and general properties

Part of managing a BIG-IP system involves configuring and maintaining a certain set of system properties. These properties fall into two main categories:

- General platform properties such as the BIG-IP system host name, IP address, and passwords for its system administrative accounts
- Device and local-traffic properties, such as NTP, DNS, and persistence settings

When you configure platform and device-related properties, you are affecting the operation of the BIG-IP system as a whole, rather than just one aspect of it. Similarly, when you configure the properties related to local traffic, you are globally affecting the operation of the local traffic management system.

◆ **Note**

*For detailed information on configuring specific features of local traffic management, see the **Configuration Guide for Local Traffic Management**.*

The remainder of this chapter describes how to configure and maintain these platform and general properties so that you can tailor the BIG-IP system to fit your needs exactly.

Configuring platform properties

When you configure platform properties, you configure settings such as the IP address of the management port, the host name of the BIG-IP system, the host IP address, and user account passwords.

You can also view information about the web server certificate, as well as import or export the certificate.

Configuring platform properties and user administration settings

From the General screen, you can configure general platform properties and user administration settings. Note that you can also configure many of these properties and settings by running the Setup utility.

Configuring general platform properties

You can configure these general properties for the BIG-IP system platform:

- An IP address, netmask, and route for the management port
- The host name for the BIG-IP system
- The host IP address for the BIG-IP system
- Whether the BIG-IP system is a single device or part of a redundant system
- The unit ID, if the system is part of a redundant system
- The time zone in which the BIG-IP system operates

The following procedure provides the basic steps for configuring platform-related general properties. Following the procedure is a description of each property, along with additional details you might need for completing step 4 of the procedure.

To configure general platform properties

1. On the Main tab of the navigation pane, expand **System**, and click **Platform**.
The General screen opens.
2. For the **Management Port** setting, type an IP address, a netmask, and a route address.
3. In the **Host Name** box, type a unique name for the BIG-IP system.
4. Configure all other general property settings as needed.
For more information, see the sections following this procedure, as well as the online help.
5. At the bottom of the screen, click **Update**.

Configuring the management port

Every BIG-IP system has a management port, also known as an administrative port. The **management port** is an interface that you use exclusively for remotely managing a BIG-IP system. You cannot use the management port for normal traffic that is slated for load balancing.

Some typical tasks that you might want to perform using the management port are:

- PXE installation and boot
- Remote management (if possible), using SSH and HTTPS
- BIG-IP system software upgrades

Configuring the management port of a BIG-IP system means assigning an IP address to the port, supplying a netmask for the management port IP address, and specifying an IP address for the BIG-IP system to use as a default route. Note that specifying a default route is only necessary if you intend to manage the BIG-IP system from a node on a different subnetwork.

◆ Note

The IP address for the management port must be in IPv4 format.

You use the **Management Port** setting on the General screen to configure the management port. There are no default values for this setting.

◆ Tip

You can also configure the management port using the LCD menu on the IP switch hardware. If you configure the management port using the LCD menu, you do not need to configure the port with the Configuration utility.

Supplying a host name

Every BIG-IP system must have a host name. Using the **Host Name** setting, type a fully qualified domain name for the BIG-IP system. An example of a host name is **mybigip.win.net**.

Assigning a host IP address

Every BIG-IP system must have a host IP address. This IP address can be the same as the address that you used for the management port, or you can assign a unique address.

To assign the host IP address, locate the **Host IP address** setting and select either **Use Management Port IP Address** or **Custom Host IP address**. The default value is **Use Management Port IP Address**.

Configuring high availability

You can use the general properties screen to specify whether the BIG-IP system is to operate as a single device or as part of a redundant system. The default value is **Single Device**.

To designate the BIG-IP system as being part of a redundant system, use the **High Availability** setting to select **Redundant Pair**. Then use the **Unit ID** setting to select the unit ID that you want to assign to the BIG-IP system (1 or 2).

Specifying a time zone

Another of the general platform properties that you can specify is the time zone. The many time zones that you can choose from are grouped into these categories: Africa, America, Antarctica, Arctic, Asia, Atlantic, Australia, Europe, Indian, and Pacific.

To set the time zone, use the **Time Zone** setting to select a time zone from the list. Select the time zone that most closely represents the location of the BIG-IP system you are configuring.

Specifying user administration settings

Part of managing platform-related properties is maintaining passwords for the system accounts, as well as enabling the **support** account. You can also configure the system to allow certain IP addresses to access the BIG-IP system through SSH.

Changing administrative account passwords

When you ran the Setup utility on the BIG-IP system, you set up some administrative accounts. Specifically, you set up the **root**, **admin**, and **support** accounts. The **root** and **admin** accounts are for use by BIG-IP system administrators, while the **support** account is for F5 Networks support personnel who require access to the customer's system for troubleshooting purposes.

Users logging in with the **root** account have console-only access to the BIG-IP system. Users logging in with the **admin** account have browser-only access to the BIG-IP system.

You can use the General screen of the platform properties to change the passwords for **root** and **admin** accounts on a regular basis. To change a password, locate the **Root Account** or **Admin Account** setting, and in the **Password** box, type a new password. In the **Confirm** box, retype the same password. For more information, see *To configure general platform properties*, on page 4-2.

Enabling the Support account

The **support** account is an optional account that you can enable on the BIG-IP system. When you enable this account, authorized F5 Networks support personnel can access the BIG-IP system to perform troubleshooting.

To enable the **support** account, find the **Support Account** setting and select **Enabled**. Then type a password, once in the **Password** box and again in the **Confirm** box.

Configuring SSH access

When you configure SSH access, you enable user access to the BIG-IP system through SSH. Also, only the IP addresses that you specify are allowed access to the system using SSH.

To configure SSH access, locate the **SSH Access** setting and click the **Enabled** box. Then use the **SSH IP Allow** setting to select either *** All Addresses** or **Specify Range**, which allows you to specify a range of addresses.

Managing a web server certificate

In addition to configuring general platform properties, you can set up a web server certificate for the BIG-IP system to use when you are managing the system through the management port. This certificate is separate from the SSL certificates that you install for managing (that is, terminating and initiating) local SSL traffic. For information on requesting and installing certificates to manage local SSL traffic, see the *Configuration Guide for Local Traffic Management*.

With respect to managing SSL for the management port, you can view information about a web server certificate that is currently installed on the BIG-IP system. You can also export a certificate or import a different certificate.

Viewing certificate and key information

You can use the Configuration utility to view information about any SSL certificate and key that you have installed on the BIG-IP system. The specific information you can view about a certificate is:

- Name
- Subject
- Expiration date
- Version
- Serial number (if any)
- Common Name
- Division
- Locality, state or province, and country
- Issuer

To view SSL certificate and key information

1. On the Main tab of the navigation pane, expand **System**, and click **Platform**.
The General screen opens.
2. On the menu bar, click **Web Server Certificate**.
This displays the properties of the certificate.

3. On the menu bar, click **Key**.
This displays the type and size of the key.

Importing, exporting, or renewing a web server certificate

You can import, export, or renew an SSL certificate.

To import an SSL certificate

1. On the Main tab of the navigation pane, expand **System** and click **Platform**.
The General screen opens.
2. On the menu bar, click **Web Server Certificate**.
This displays the properties of the certificate.
3. At the bottom of the screen, click **Import**.
The Web Server Certificate screen opens.
4. From the **Import Type** list, select an import type, either **Certificate** or **Certificate and Key**.
5. From the **Certificate Source** setting, click either **Upload File** or **Paste Text**:
 - If you click **Upload File**, type a file name or click **Browse**.
If you click **Browse**:
 - a) Navigate to the relevant Windows folder and click a file name.
 - b) On the browser window, click **Open**.
 - If you click **Paste Text**:
 - a) Copy the text from another source.
 - b) Paste the text into the **Certificate Source** window.
6. Click **Import**.

To export an SSL certificate

1. On the Main tab of the navigation pane, expand **System**, and click **Platform**.
The General screen opens.
2. On the menu bar, click **Web Server Certificate**.
This displays the properties of the certificate.
3. At the bottom of the screen, click **Export**.
The screen displays the text of the existing certificate.
4. Next to the **Certificate File** setting, click **Download server.crt**.

To renew an SSL certificate

1. On the Main tab of the navigation pane, expand **System**, and click **Platform**.
The General screen opens.
2. On the menu bar, click **Web Server Certificate**.
This displays the properties of the certificate.
3. At the bottom of the screen, click **Renew**.
This displays the properties of the certificate and its associated key.
4. Change any properties as needed.
For detailed information, see the online help.
5. Click **Finished**.

Importing and exporting a key

You can import and export SSL keys.

To import an SSL key

1. On the Main tab of the navigation pane, expand **System**, and click **Platform**.
The General screen opens.
2. On the menu bar, click **Web Server Certificate**.
This displays the properties of the certificate.
3. On the menu bar, click **Key**.
4. Click **Import**.
5. From the **Import Type** list, select an import type, either **Certificate** or **Certificate and Key**.
6. From the **Key Source** setting, click either **Upload File** or **Paste Text**:
 - If you click **Upload File**, type a file name or click **Browse**.
If you click **Browse**:
 - a) Navigate to the relevant Windows folder and click a file name.
 - b) On the browser window, click **Open**.
 - If you click **Paste Text**:
 - a) Copy the text from another source.
 - b) Paste the text into the **Key Source** window.
7. Click **Import**.

To export an SSL key

1. On the Main tab of the navigation pane, expand **System**, and click **Platform**.
The General screen opens.
2. On the menu bar, click **Web Server Certificate**.
This displays the properties of the certificate.
3. On the menu bar, click **Key**.
This displays the properties for that key.
4. Click **Export**.
The screen displays the text of the key.
5. Next to the **Key File** setting, click **Download server.key**.

Configuring general properties

Using the Configuration utility, you can view and configure a number of general BIG-IP system properties. Some of these properties are related to the BIG-IP system as a device, while others are related to local traffic management.

Configuring device-related properties

You can view or configure a number of properties related to the BIG-IP system as a device. These properties fall into three main categories: general device properties, Network Time Protocol (NTP) properties, and Domain Name System (DNS) properties.

Configuring general properties

The BIG-IP general properties that you can view are:

- The host name
- The BIG-IP software version number
- The number of CPUs available
- The number of CPUs that are active
- The current CPU mode (uniprocessor or multiprocessor)

The BIG-IP general properties that you can configure are:

- Network boot
- Quiet boot
- The percent of memory usage for reboot

The following procedure provides the basic steps for configuring general properties. Following the procedure are descriptions of the properties that you might need for completing step 2 of the procedure.

To view or configure general properties

1. On the Main tab of the navigation pane, expand **System**, and click **General Properties**.
The General screen opens.
2. View or configure any settings.
For detailed information on these settings, see the online help and Table 4.1, on page 4-10.
3. If you configured any settings, click **Update**.

Table 4.1, on page 4-10 lists and describes the general properties that you can view or configure.

Property	Description	Default Value
Host Name	Displays the host name of the BIG-IP system. This name is the same host name that you specified on the main Platform screen.	No default value
Version	Displays the version number of the BIG-IP system software that is running on the system.	No default value
CPU Count	Displays the total number of CPUs that the BIG-IP system contains.	No default value
Active CPUs	Displays the total number of CPUs that are currently active on the BIG-IP system.	No default value
CPU Mode	Displays the current processor mode of the system, either uniprocessor or multiprocessor.	No default value
Network Boot	Enables or disables the network boot feature. If you enable this feature and then reboot the system, the system boots from an ISO image on the network, rather than from an internal media drive. Use this option only when you want to install software on the system, for example, for an upgrade or a re-installation. Note that this setting reverts to Disabled after you reboot the system a second time.	Disabled (unchecked)
Quiet Boot	Enables or disables the quiet boot feature. If you enable this feature, the system suppresses informational text on the console during the boot cycle.	Enabled (checked)
Memory Restart Percent	Specifies the memory usage percent at which the system reboots.	97

Table 4.1 General properties of a BIG-IP system device

Configuring NTP

Network Time Protocol (NTP) is a protocol that synchronizes the clocks on a network. You can use the Configuration utility to specify a list of IP addresses of the servers that you want the BIG-IP system to use when updating the time on network systems. You can also edit or delete the entries in the server list.

To configure a list of NTP time servers

1. On the Main tab of the navigation pane, expand **System**, and click **General Properties**.
The General screen opens.
2. From the Device menu, choose NTP.
This opens the NTP screen.
3. For the **Time Server List** setting, add, edit, or remove an IP address:

- To add an IP address to the list:
 - a) Type a time server's IP address or host name in the **Address** box.
 - b) Click **Add**.
 - To edit an IP address in the list:
 - a) In the **Time Server List** area, select an IP address.
The IP address appears in the **Address** box.
 - b) In the **Address** box, change the IP address.
 - c) Click the **Edit** button.
 - To remove an IP address from the list:
 - a) In the **Time Server List** area, select an IP address.
The IP address appears in the **Address** box.
 - b) Click the **Delete** button.
4. Click **Update**.

Configuring DNS

Domain Name System (DNS) is an industry-standard distributed internet directory service that resolves domain names to IP addresses. If you plan to use DNS in your network, you can use the Configuration utility to configure DNS for the BIG-IP system.

When you configure DNS, you create two lists: a DNS lookup server list, and a BIND forwarder server list. The *DNS lookup server list* allows BIG-IP system users to use IP addresses, host names, or fully-qualified domain names (FQDNs) to access virtual servers, nodes, or other network objects.

The *BIND forwarder server list* provides DNS resolution for servers and other equipment load balanced by the BIG-IP system, that is, for the servers that the BIG-IP system uses for DNS proxy services.

◆ Note

*To use DNS Proxy services, you must enable the **named** service.*

In addition to adding servers to the DNS lookup server list and the BIND forwarder server list, you can also edit or delete the entries in these lists.

To configure DNS for the BIG-IP system

1. On the Main tab of the navigation pane, expand **System**, and click **General Properties**.
The General screen opens.
2. From the Device menu, choose DNS.
This opens the DNS screen.

3. In the **DNS Lookup Server List** area, you can add, edit, or remove a server IP address:
 - To add a server to the list:
 - a) Type the IP address of a properly-configured name server in the **Address** box.
 - b) Click **Add**.
 - c) To add backup DNS servers to the list, repeat steps a and b.
 - To edit an IP address in the list:
 - a) In the **DNS Lookup Server List** area, select an IP address. The IP address appears in the **Address** box.
 - b) In the **Address** box, change the IP address.
 - c) Click **Edit**.
 - To remove an IP address from the list:
 - a) In the **DNS Lookup Server List** area, select an IP address. The IP address appears in the **Address** box
 - b) Click **Delete**.
4. In the **BIND Forwarder Server List** area, you can add, edit, or remove a server IP address:
 - To add a server to the list:
 - a) Type a server's IP address in the **Address** box.
 - b) Click **Add**.
 - To edit an IP address in the list:
 - a) In the **BIND Forwarder Server List** area, select an IP address. The IP address appears in the **Address** box.
 - b) In the **Address** box, change the IP address.
 - c) Click **Edit**.
 - To remove an IP address from the list:
 - a) In the **BIND Forwarder Server List** area, select an IP address. The IP address appears in the **Address** box
 - b) Click **Delete**.
5. Click **Update**.

Configuring local-traffic properties

The BIG-IP system includes a set of properties that apply globally to the local traffic management system. These properties fall into two main categories: general local-traffic properties, and persistence properties. You can use the Configuration utility to configure and maintain these properties.

Configuring general local-traffic properties

You can configure a number of properties that affect the general behavior of the BIG-IP local traffic management system. In most cases, these properties are not directly related to any one type of local traffic management object, such as a virtual server or a load balancing pool.

The following procedure provides the basic steps for configuring general local-traffic properties. Following the procedure are descriptions of the properties with additional details you might need for completing step 3 of the procedure.

To configure general local-traffic properties

1. On the Main tab of the navigation pane, expand **System**, and click **General Properties**.
The General screen opens.
2. From the Local Traffic menu, choose General.
3. Configure all properties or retain the default values.
For detailed information, see Table 4.2.
4. Click **Update**.

Table 4.2 lists and describes the properties that you can configure to manage the behavior of the local traffic management system.

Property	Description	Default Value
Auto Last Hop	Specifies, when checked (enabled), that the system automatically maps the last hop for pools.	Enabled (checked)
Maintenance Mode	Specifies, when checked (enabled), that the unit is in maintenance mode. In maintenance mode, the system stops accepting new connections and slowly completes the processing of existing connections.	Disabled (unchecked)
VLAN-Keyed Connections	Check this setting to enable VLAN-keyed connections. VLAN-keyed connections are used when traffic for the same connection must pass through the system several times, on multiple pairs of VLANs (or in different VLAN groups).	Enabled (checked)

Table 4.2 General properties for globally managing local traffic

Property	Description	Default Value
Path MTU Discovery	Specifies, when checked (enabled), that the system discovers the maximum transmission unit (MTU) that it can send over a path without fragmenting TCP packets.	Enabled (checked)
Reject Unmatched Packets	Specifies, when checked (enabled), that the system returns a TCP RESET or ICMP_UNREACH packet if no virtual servers on the system match the destination address of the incoming packet. When this setting is disabled, the system silently drops the unmatched packet.	Enabled (checked)
Maximum Node Idle Time	Specifies the number of seconds a node can be left idle by the Fastest load balancing mode. The system sends fewer connections to a node that is responding slowly, and periodically recalculates the response time of the slow node.	0 (disabled)
Reaper High-water Mark	Specifies, in percent, the memory usage at which the system stops establishing new connections. Once the system meets the reaper high-water mark, the system does not establish new connections until the memory usage drops below the reaper low-water mark. To disable the adaptive reaper, set the high-water mark to 100 . <i>Note: This setting helps to mitigate the effects of a denial-of-service attack.</i>	95
Reaper Low-water Mark	Specifies, in percent, the memory usage at which the system silently purges stale connections, without sending reset packets (RST) to the client. If the memory usage remains above the low-water mark after the purge, then the system starts purging established connections closest to their service timeout. To disable the adaptive reaper, set the low-water mark to 100 .	85
SYN Check™ Activation Threshold	Specifies the number of new or untrusted TCP connections that can be established before the system activates the SYN Cookies authentication method for subsequent TCP connections.	16384
Layer 2 Cache Aging Time	Specifies, in seconds, the amount of time that records remain in the Layer 2 forwarding table, when the MAC address of the record is no longer detected on the network.	300

Table 4.2 General properties for globally managing local traffic

Property	Description	Default Value
Share Single MAC Address	Specifies, when checked (enabled), that all VLANs share a single MAC address. If you use the default value (unchecked), the BIG-IP gives each VLAN the MAC address of the VLAN's lowest-numbered interface. Use this setting when configuring an active/standby redundant system.	Disabled (unchecked)
SNAT Packet Forwarding	Specifies the type of traffic for which the system attempts to forward (instead of reject) Any-IP packets, when the traffic originates from a member of a SNAT. There are two possible values: TCP and UDP Only: Specifies that the system forwards, for TCP and UDP traffic only, Any-IP packets originating from a SNAT member. All Traffic: Specifies that the system forwards, for all traffic types, Any-IP packets originating from a SNAT member.	TCP and UDP Only

Table 4.2 General properties for globally managing local traffic

Configuring persistence properties

Using the Configuration utility, you can perform certain persistence-related tasks such as managing the way that destination IP addresses are stored in the persistence table, and specifying a data group that contains proxy IP addresses.

The following procedure provides the basic steps for configuring general persistence-related properties. Following the procedure are descriptions of the properties with additional details you might need for completing step 3 of the procedure.

To configure persistence properties

1. On the Main tab of the navigation pane, expand **System**, and click **General Properties**.
The General screen opens.
2. From the Local Traffic menu, choose Persistence.
3. Configure the properties or retain the default values.
For detailed information, see Table 4.3, on page 4-16.
4. Click **Update**.

Table 4.3 lists and describes the properties that you can configure to manage general persistence-related properties.

Property	Description	Default Value
Management of Destination Address Entries	<p>Specifies how the system manages the destination IP address entries in the persistence table.</p> <p>Timeout: Specifies that entries remain in the persistence table until the BIG-IP system times them out, based on the timeout value configured in the corresponding persistence profile.</p> <p>Maximum Entries: Specifies that the system stops adding entries to the persistence table when the number of entries reaches the maximum number of entries allowed.</p>	Timeout
Proxy Address Data Group	<p>Specifies the data group that contains proxy IP addresses. You use this data group to identify the addresses that are to be treated as proxies when you enable the Map Proxies option on a persistence profile.</p>	aol

Table 4.3 General properties for controlling session persistence



5

Working with Interfaces

- Introducing BIG-IP system interfaces
- Configuring interfaces
- Configuring interface mirroring
- Displaying interface statistics
- Related configuration tasks

Introducing BIG-IP system interfaces

A key task of the BIG-IP system configuration is the configuration of BIG-IP system interfaces. The *interfaces* on a BIG-IP system are the physical ports that you use to connect the BIG-IP system to other devices on the network. These other devices can be next-hop routers, layer 2 devices, destination servers, and so on. Through its interfaces, the BIG-IP system can forward traffic to or from other network devices.

◆ **Note**

*Throughout this guide, the term **interface** refers to the physical ports on the BIG-IP system.*

Every BIG-IP system includes multiple interfaces. The exact number of interfaces that you have on the BIG-IP system depends on the platform type. For information on BIG-IP platform types, see *Platform Guide: 1500, 3400, 6400, and 6800*.

Each of the interfaces on the BIG-IP system has unique properties, such as media speed, duplex mode, VLAN tagging, and spanning tree protocol settings. You can use the Configuration utility to configure these properties. For more information, see *Configuring interfaces*, on page 5-2.

In addition to configuring interface properties, you can implement a feature known as interface mirroring, which you can use to duplicate traffic from one or more interfaces to another. You can also view statistics about the traffic on each interface. For more information, see *Configuring interface mirroring*, on page 5-6 and *Displaying interface statistics*, on page 5-7.

Once you have configured the properties of each interface, you can configure several other features of the BIG-IP system that control the way that interfaces operate. For example, by creating a virtual local area network (VLAN) and assigning interfaces to it, the BIG-IP system can insert a VLAN ID, or *tag*, into frames passing through those interfaces. In this way, a single interface can forward traffic for multiple VLANs. For more information on configuring other BIG-IP features related to interfaces, see *Related configuration tasks*, on page 5-9.

Configuring interfaces

Each interface on the BIG-IP system has a set of properties that you can configure, such as enabling or disabling the interface, setting the requested media type and duplex mode, and configuring flow control. Configuring the properties of each interface is one of the first tasks you do after running the Setup utility on the BIG-IP system. While you can change some of these properties, such as media speed and duplex mode, you cannot change other properties, such as the media access control (MAC) address.

One of the interfaces on the BIG-IP system is a special interface dedicated to performing a specific set of system management functions. This interface is called the *management interface*, named **MGMT**. For information on how to configure and use the management interface, see Chapter 4, *Configuring the BIG-IP Platform and General Properties*.

◆ **Note**

*For information on configuring STP-related properties on an interface, see Chapter 9, **Configuring Spanning Tree Protocols**.*

Before configuring interface properties, it is helpful to understand interface naming conventions.

Understanding interface naming conventions

By convention, the names of the interfaces on the BIG-IP system use the format **<s>.<p>** where **s** is the slot number of the network interface card (NIC), and **p** is the port number on the NIC. Examples of interface names are 1.1, 1.2, and 2.1. BIG-IP system interfaces already have names assigned to them; you do not explicitly assign them.

An exception to the interface naming convention is the management interface, which has the special name **MGMT**. For more information on the management interface, see Chapter 4, *Configuring the BIG-IP Platform and General Properties*.

Viewing interface information and media properties

Using the Configuration utility, you can display a screen that lists all of the BIG-IP system interfaces, as well as their current status (**UP** or **DOWN**). You can also view other information about each interface:

- The MAC address of the interface
- The media type (such as 100baseTX or 100base FX)
- The number of VLANs to which the interface is assigned
- A trunk name, if the interface is assigned to a trunk

This information is useful when you want to assess the way that a particular interface is forwarding traffic. For example, you can use this information to determine the specific VLANs for which an interface is currently forwarding traffic. You can also use this information to determine the speed at which an interface is currently operating.

On the General Properties screen for interfaces, you can view the media speed and the duplex mode of an interface. Use the following procedures to view the list of interfaces and related information, and to view the media properties of an interface.

To view a list of interfaces and related information

On the Main tab of the navigation pane, expand **Network**, and click **Interfaces**. This displays a list of the interfaces on the BIG-IP system, along with their status and related information.

To view media properties for an interface

1. On the Main tab of the navigation pane, expand **Network**, and click **Interfaces**.
This displays the list of the interfaces on the BIG-IP system.
2. Click an interface name in the list.
This displays the general properties of that interface, as well as some configuration settings.
3. In the General Properties area of the screen, view the **Media Speed** and **Active Duplex** properties.

Configuring interface properties

You can configure a number of general properties for each interface. When you configure these properties, you customize the way that the interface forwards traffic. For example, if you want the interface to operate as part of a trunk, you can set the **Requested Duplex** mode to **full**, which is a requirement for trunk participation. Table 5.1 lists and describes these properties.

Interface Properties	Description	Default Value
State	Enables or disables the interface.	Enabled
Requested media	Specifies a media type or auto for automatic detection.	auto
Requested Duplex	Specifies that the interface is in full or half duplex mode.	half
Flow Control	Specifies how an interface handles pause frames for flow control.	Pause TX/RX

Table 5.1 The properties you can configure for an interface

Use the following procedure to configure the general properties of an interface. For detailed information on these individual properties, see the sections following the procedure.

◆ **Note**

*To configure STP-related settings, see Chapter 9, **Configuring Spanning Tree Protocols**.*

To configure general interface properties

1. On the Main tab of the navigation pane, expand **Network**, and click **Interfaces**.
This displays the list of the interfaces on the BIG-IP system.
2. In the Name column, click an interface name.
This displays the general properties of that interface, as well as some configuration settings.
3. In the Configuration area, configure the properties as needed.
For information on each property, see the following sections.
4. Click **Update**.

Configuring the state of an interface

You can either enable or disable an interface on the BIG-IP system, by configuring the **State** property. By default, each interface is set to **Enabled**, where it can accept ingress or egress traffic. When you set the state to **Disabled**, the interface cannot accept ingress or egress traffic.

Setting the requested media type

You can configure the **Requested Media** property to specify the media type of the interface card, or you can use the **auto** setting for auto-detection. The default setting for this property is **auto**. If the media type is set to **auto** and the card does not support auto-detection, the default type for that interface is used, for example **1000BaseTX**.

The values that you can choose from when configuring the **Requested Media** property are: **none**, **auto**, **10baseT**, **100baseTX**, and **1000baseTX**.

◆ **Note**

If the BIG-IP system is inter-operating with an external switch, the media setting should match that of the external switch.

Setting the duplex mode

You can configure the **Requested Duplex** property to set the duplex mode to full or half duplex. When the **Media Type** property is set to **auto**, the default setting for the **Requested Duplex** property is **full**. **Full duplex** mode means that traffic on that interface can travel in both directions

simultaneously, while *half duplex* mode means that traffic on that interface can only travel in one direction at any given time. Note that if you want the interface to be part of a trunk, it must be set to full duplex mode.

If the media type of the interface does not allow the duplex mode to be set, this is indicated by an on-screen message. If setting the duplex mode is not supported for the interface, the duplex setting is not saved to the **bigip_base.conf** file.

The values that you can choose from when configuring the **Requested Duplex** property are **half** and **full**.

Configuring flow control

You can configure the **Flow Control** property to manage the way that an interface handles pause frames for flow control. *Pause frames* are frames that an interface sends to a peer interface as a way to control frame transmission from that peer interface. Pausing a peer's frame transmissions prevents an interface's First-in, First-out (FIFO) queue from filling up and resulting in a loss of data. Possible values for this property are:

- **Pause None**
Disables flow control.
- **Pause TX/RX**
Specifies that the interface honors pause frames from its peer, and also generates pause frames when necessary. This is the default value.
- **Pause TX**
Specifies that the interface ignores pause frames from its peer, and generates pause frames when necessary.
- **Pause RX**
Specifies that the interface honors pause frames from its peer, but does not generate pause frames.

Configuring interface mirroring

For reliability reasons, you can configure a feature known as interface mirroring. When you configure *interface mirroring*, you cause the BIG-IP system to copy the traffic on one or more interfaces to another interface that you specify. By default, the interface mirroring feature is disabled.

The settings you configure to implement the interface mirroring feature are shown in Table 5.2.

Setting	Description
Interface Mirroring State	Enables or disables interface mirroring.
Destination Interface	Specifies the interface on which traffic from other interfaces is to be mirrored.
Mirrored Interfaces	Specifies one or more interfaces for which you want to mirror traffic on the destination interface.

Table 5.2 Configuration settings for enabling interface mirroring

Use the following procedure to configure interface mirroring.

To configure interface mirroring

1. On the Main tab of the navigation pane, expand **Network**, and click **Interfaces**.
This displays the list of the interfaces on the BIG-IP system.
2. On the menu bar, click **Interface Mirroring**.
The Interface Mirroring screen opens.
3. From the **Interface Mirroring State** list, select Enabled.
This displays additional configuration settings.
4. From the **Destination Interface** list, select the interface that you want the BIG-IP system to use for mirrored traffic.
5. For the **Mirrored Interfaces** setting, click an interface number in the **Available** box, and using the Move button (<<), move the interface number to the **Selected** box. Repeat this step for each interface that you want to mirror.
6. Click **Update**.

Displaying interface statistics

You can display a variety of statistics about the interfaces on the BIG-IP system. Figure 5.1 shows an example of the output you see when you display interface statistics on an active unit of a redundant system.

Interface Statistics			Bits		Packets		Multicast		Errors		Drops		
<input type="checkbox"/>	▲ Name	Status	In	Out	In	Out	In	Out	In	Out	In	Out	Collisions
<input type="checkbox"/>	1.1	UP	63.3M	23.8M	24.0K	32.7K	0	17	0	0	0	0	0
<input type="checkbox"/>	1.2	DOWN	0	0	0	0	0	0	0	0	0	0	0
<input type="checkbox"/>	1.3	DOWN	0	0	0	0	0	0	0	0	0	0	0
<input type="checkbox"/>	1.4	DOWN	0	0	0	0	0	0	0	0	0	0	0
<input type="checkbox"/>	2.1	UNPOPULATED	0	0	0	0	0	0	0	0	0	0	0
<input type="checkbox"/>	2.2	UNPOPULATED	0	0	0	0	0	0	0	0	0	0	0

Figure 5.1 Sample interface statistics screen

◆ Tip

For descriptions of each type of statistic, see the online help.

To display interface statistics

1. On the Main tab of the navigation pane, expand **Network**, and click **Interfaces**.
This displays the list of the interfaces on the BIG-IP system.
2. On the menu bar, click **Statistics**.
This displays statistics for the interfaces on the BIG-IP system.
3. In the Display Options area of the screen, verify that the **Statistics Type** setting is set to **Interfaces**.

4. For the **Data Format** setting, retain the default value (**Normalized**), or select **Unformatted** from the list.
5. For the **Auto Refresh** setting, retain the default value of **Disabled**, or select an automatic refresh interval from the list.

*Note: Setting the **Auto Refresh** value to a short interval could impact system performance.*

6. If you want to manually refresh the statistics, click **Refresh**.

Related configuration tasks

After you have configured the interfaces on the BIG-IP system, one of the primary tasks you perform is to assign those interfaces to the virtual LANs (VLANs) that you create. A **VLAN** is a logical subset of hosts on a local area network (LAN) that reside in the same IP address space. When you assign multiple interfaces to a single VLAN, traffic destined for a host in that VLAN can travel through any one of these interfaces to reach its destination. Conversely, when you assign a single interface to multiple VLANs, the BIG-IP system can use that single interface for any traffic that is intended for hosts in those VLANs. For more information on VLANs and assigning interfaces to them, see Chapter 7, *Configuring VLANs and VLAN Groups*.

Another powerful feature that you can use for BIG-IP system interfaces is trunking, with link aggregation. A **trunk** is an object that logically groups physical interfaces together to increase bandwidth. Link aggregation, through the use of the industry-standard Link Aggregation Control Protocol (LACP), provides regular monitoring of link status, as well as failover if an interface becomes unavailable. For more information on using trunks and LACP, see Chapter 8, *Working with Trunks*.

Finally, you can configure your BIG-IP system interfaces to work with one of the spanning tree protocols (STP, RSTP, and MSTP). **Spanning tree protocols** reduce traffic on your internal network by blocking duplicate routes to prevent bridging loops. Chapter 9, *Configuring Spanning Tree Protocols*, describes the spanning tree protocols and the procedure for configuring these protocols on the BIG-IP system. The chapter also includes information on setting spanning tree-related properties on individual interfaces.



6

Configuring Self IP Addresses

- Introducing self IP addresses
- Creating and managing self IP addresses

Introducing self IP addresses

A *self IP address* is an IP address that you associate with a VLAN, to access hosts in that VLAN. By virtue of its netmask, a self IP address represents an *address space*, that is, a range of IP addresses spanning the hosts in the VLAN, rather than a single host address. You can associate self IP addresses not only with VLANs, but also with VLAN groups.

Self IP addresses serve two purposes. First, when sending a message to a destination server, the BIG-IP system uses the self IP addresses of its VLANs to determine the specific VLAN in which a destination server resides. For example, if VLAN **internal** has a self IP address of **10.10.10.100**, with a netmask of **255.255.255.0**, and the destination server's IP address is **10.10.10.20** (with a netmask of **255.255.255.255**), the BIG-IP system recognizes that the server's IP address falls within the range of VLAN **internal**'s self IP address, and therefore sends the message to that VLAN. More specifically, the BIG-IP system sends the message to the interface that you assigned to that VLAN. If more than one interface is assigned to the VLAN, the BIG-IP system takes additional steps to determine the correct interface, such as checking the layer2 forwarding table.

Second, a self IP address serves as the default route for each destination server in the corresponding VLAN. In this case, the self IP address of a VLAN appears as the destination IP address in the packet header when the server sends a response to the BIG-IP system. For more information on configuring the default route of a destination server, see Chapter 12, *Setting up a Redundant System*.

You normally assign self IP addresses to a VLAN when you initially run the Setup utility on a BIG-IP system. More specifically, you assign one static self IP address and one floating self IP address to each of the default VLANs (**internal** and **external**). Later, using the Configuration utility, you can create self IP addresses for other VLANs that you create.

Types of self IP addresses

There are two types of self IP addresses that you can create:

- A *static self IP address* is an IP address that the BIG-IP system does not share with another BIG-IP system. By default, the self IP addresses that you create with the Configuration utility are static self IP addresses.
- A *floating self IP address* is an IP address that two BIG-IP systems share, such as two units of a redundant system. When you use the Configuration utility to create a self IP address, you can specify that you want the IP address to be floating address.

For more information on static and floating IP addresses, see Chapter 12, *Setting up a Redundant System*.

Self IP addresses and MAC addresses

For each self IP address that you create for a VLAN, the BIG-IP system automatically assigns a media access control (MAC) address. By default, the BIG-IP system assigns the same MAC address that is assigned to the lowest-numbered interface of the VLAN.

As an alternative, you can globally configure the BIG-IP system to assign the same MAC address to all VLANs. This feature is useful if your network includes a type of switch that does not keep a separate layer 2 forwarding table for each VLAN on that switch.

Using self IP addresses for SNATs

When you configure the BIG-IP system to manage local area traffic, you can implement a feature known as a secure network address translation (SNAT). A *SNAT* is an object that causes the BIG-IP system to translate the original source IP address of a packet to an IP address that you specify. A SNAT ensures that the target server sends its response back through the BIG-IP system rather than to the original client IP address directly.

When you create a SNAT, you can configure the BIG-IP system to automatically choose a translation address. This ability of the BIG-IP system to automatically choose a translation address is known as *SNAT automapping*, and in this case, the translation address that the system chooses is always an existing self IP address. Thus, for traffic going from the BIG-IP system to a destination server, configuring SNAT automapping ensures that the source IP address in the header of a packet is a self IP address.

When you create an automapped SNAT, the BIG-IP system actually creates a SNAT pool consisting of the system's internal self IP addresses, and then uses an algorithm to select and assign an address from that SNAT pool.

For more information on SNAT automapping, see the *Configuration Guide for Local Traffic Management*.

Creating and managing self IP addresses

As stated previously, it is when you initially run the Setup utility on a BIG-IP system that you normally create any static and floating self IP addresses and assign them to VLANs. However, if you want to create additional self IP addresses later, you can do so using the Configuration utility.

Creating a self IP address

The BIG-IP system offers several settings that you can configure for a self IP address. These settings are summarized in Table 6.1.

Setting	Description	Default Value
IP Address	Specifies a self IP address.	No default value
Netmask	Specifies the netmask for the self IP address.	No default value
VLAN	Specifies the VLAN to which this self IP address corresponds.	No default value
Port Lockdown	Specifies the protocols and services from which the self IP address can accept traffic.	Allow Default
Floating IP	Specifies that the self IP address is a floating IP address (shared between two BIG-IP systems).	Unchecked

Table 6.1 Configuration settings for a self IP address

Use the following procedure to create a self IP address. For detailed information about each setting, see the sections following the procedure.

◆ Note

A self IP address can be in either IPv4 or IPv6 format.

To create a self IP address

1. On the Main tab of the navigation pane, expand **Network**, and click **Self IPs**.
This displays a list of existing self IP addresses.
2. In the upper-right corner of the screen, click the **Create** button.
3. In the **IP Address** box, type the self IP address that you want to assign to a VLAN.
4. In the **Netmask** box, type a netmask.
5. For the **VLAN** setting, select the name of the VLAN to which you want to assign the self IP address.
The default value is **internal**.

6. For the **Port Lockdown** setting, select **Allow Default**, **Allow All**, **Allow None**, or **Allow Custom**.
Selecting **Allow Custom** displays the **Custom List** setting. For more information on these setting values, see *Specifying port lockdown*, on page 6-5.
7. If you chose **Allow Custom** in step 7, click **TCP**, **UDP**, or **Protocol**.
 - a) If you chose **TCP** or **UDP**, do one or both of the following:
 - Click **All** or **None** and then click **Add**.
The value **All** or **None** appears in the **TCP** or **UDP** box.
 - Click **Port**, type a port number, and then click **Add**.
The port number appears in the **TCP** or **UDP** box.
 - b) If you chose **Protocol**, select a protocol name and click **Add**.
8. If you want to configure the self IP address as a floating IP address, check the **Floating IP** box.
9. To finish the configuration of this self IP address and create other self IP addresses, click **Repeat** and perform all previous steps until all self IP addresses have been created.
10. Click **Finished**.

Specifying an IP address

As described in *Introducing self IP addresses*, on page 6-1, a self IP address, combined with a netmask, typically represents a range of host IP addresses in a VLAN. If you are assigning a self IP address to a VLAN group, the self IP address represents the range of self IP addresses assigned to the VLANs in that group.

The self IP address that you specify in the **IP Address** setting is a static IP address, unless you enable the **Floating IP** setting. For more information, see *Specifying a floating IP address*, on page 6-5.

Specifying a netmask

When you specify a netmask for a self IP address, the self IP address can represent a range of IP addresses, rather than a single host address. For example, a self IP address of **10.0.0.100** can represent several host IP addresses if you specify a netmask of **255.255.0.0**.

Associating a self IP address with a VLAN

You assign a unique self IP address to a specific VLAN or a VLAN group:

- ◆ **Assigning a self IP address to a VLAN**

The self IP address that you assign to a VLAN should represent an address space that includes the self IP addresses of the hosts that the VLAN contains. For example, if the address of one destination server in

a VLAN is **10.0.0.1** and the address of another server in the VLAN is **10.0.0.2**, you could assign a self IP address of **11.0.0.100**, with a netmask of **255.255.0.0**, to the VLAN.

◆ **Assigning a self IP address to a VLAN group**

The self IP address that you assign to a VLAN group should represent an address space that includes the self IP addresses of the VLANs that you assigned to the group. For example, if the self IP address of one VLAN in a VLAN group is **10.0.20.100** and the address of the other VLAN in a VLAN group is **10.0.30.100**, you could assign an address of **10.0.0.100**, with a netmask of **255.255.0.0**, to the VLAN group.

The **VLAN** list displays the names of all existing VLANs and VLAN groups.

Specifying port lockdown

Each self IP address has a feature known as port lockdown. **Port lockdown** is a security feature that allows you to specify particular UDP and TCP protocols and services from which the self IP address can accept traffic. By default, a self IP address accepts traffic from these protocols and services:

- For UDP, the allowed protocols and services are: DNS (53), SNMP (161), RIP (520)
- For TCP, the allowed protocols and services are: SSH (22), DNS (53), SNMP (161), HTTPS (443), 4353 (iQuery)

If you do not want to use the default setting (**Allow Default**), you can configure port lockdown to allow either all UDP and TCP protocols and services (**Allow All**), no UDP protocols and services (**Allow None**), or only those that you specify (**Allow Custom**).

Specifying a floating IP address

You can enable the **Floating IP** setting if you want the self IP address to be a floating IP address, that is, an address shared between two BIG-IP systems. A floating self IP address enables a destination server to successfully send a response when the relevant BIG-IP unit is unavailable. When two units share a floating self IP address, a destination server can send traffic to that address instead of a static self IP address. If the target unit is unavailable, the peer unit can receive and process that traffic. Without this shared floating IP address, the delivery of server traffic to a unit of a redundant system can fail.

◆ **Note**

*The **Floating IP** setting appears on the screen only when the BIG-IP system is configured as a unit of a redundant system. For more information on configuring a redundant system, see Chapter 12, **Setting up a Redundant System**.*

Managing self IP addresses

Using the Configuration utility, you can view or change the properties of a self IP address, or delete a self IP address.

To view or modify the settings of a self IP address

1. On the Main tab of the navigation pane, expand **Network**, and click **Self IPs**.
This displays a list of existing self IP addresses.
2. In the IP Address column, click a self IP address.
This displays the properties page for that self IP address.
3. To change the setting values, modify the values and click **Update**.
This displays the list of existing self IP addresses.

◆ Note

*You can modify any setting except **IP Address**. To modify the **IP Address** setting, you must delete the self IP address and create a new one. For more information, see **To delete a self IP address**, following, and **To create a self IP address**, on page 6-3.*

To delete a self IP address

1. On the Main tab of the navigation pane, expand **Network**, and click **Self IPs**.
This displays a list of existing self IP addresses.
2. In the IP address column, locate the self IP address you want to delete.
3. Click the select box to the left of the IP address.
4. Click **Delete**.
A confirmation screen appears.
5. Click **Delete**.



7

Configuring VLANs and VLAN Groups

- Introducing virtual LANs
- Creating and managing VLANs
- Creating and managing VLAN groups
- Assigning self IPs to VLANs and VLAN groups

Introducing virtual LANs

In Chapter 1, *Introducing BIG-IP Network and System Management*, we described the BIG-IP system as being a multilayer switch instead of a standard IP router. This allows you to create and deploy virtual local area networks (VLANs). A **VLAN** is a logical subset of hosts on a local area network (LAN) that operate in the same IP address space. Grouping hosts together in a VLAN has distinct advantages. For example, with VLANs, you can:

- Reduce the size of broadcast domains, thereby enhancing overall network performance.
- Reduce system and network maintenance tasks substantially. Functionally-related hosts no longer need to physically reside together to achieve optimal network performance.
- Enhance security on your network by segmenting hosts that must transmit sensitive data.

The way that you group hosts into VLANs is by using the Configuration utility to create a VLAN and associate physical interfaces with that VLAN. In this way, any host that sends traffic to a BIG-IP system interface is logically a member of the VLAN or VLANs to which that interface belongs.

Understanding VLANs on a BIG-IP system

The BIG-IP system is a port-based switch that includes multilayer processing capabilities. These capabilities enhance standard VLAN behavior, in these ways:

- You can associate physical interfaces on the BIG-IP system directly with VLANs. In this way, you can associate multiple interfaces with a single VLAN, or you can associate a single interface with multiple VLANs.
- You do not need physical routers to establish communication between separate VLANs. Instead, the BIG-IP system can process messages between VLANs.
- You can incorporate a BIG-IP system into existing, multi-vendor switched environments, due to the BIG-IP system's compliance with the IEEE 802.1q VLAN standard.
- You can combine two or more VLANs into an object known as a VLAN group. With a **VLAN group**, a host in one VLAN can communicate with a host in another VLAN using a combination of layer 2 forwarding and IP routing. This offers both performance and reliability benefits.

Understanding the default VLAN configuration

By default, the BIG-IP system includes two VLANs, named **internal** and **external**. When you initially ran the Setup utility, you assigned the following to each of these VLANs:

- A static and a floating self IP address
- A VLAN tag
- One or more BIG-IP system interfaces

A typical VLAN configuration is one in which you create the two VLANs **external** and **internal**, and one or more BIG-IP system interfaces assigned to each VLAN. You then create a virtual server, and associate a default load balancing pool with that virtual server. Figure 7.1 shows a typical configuration using the default VLANs **external** and **internal**.

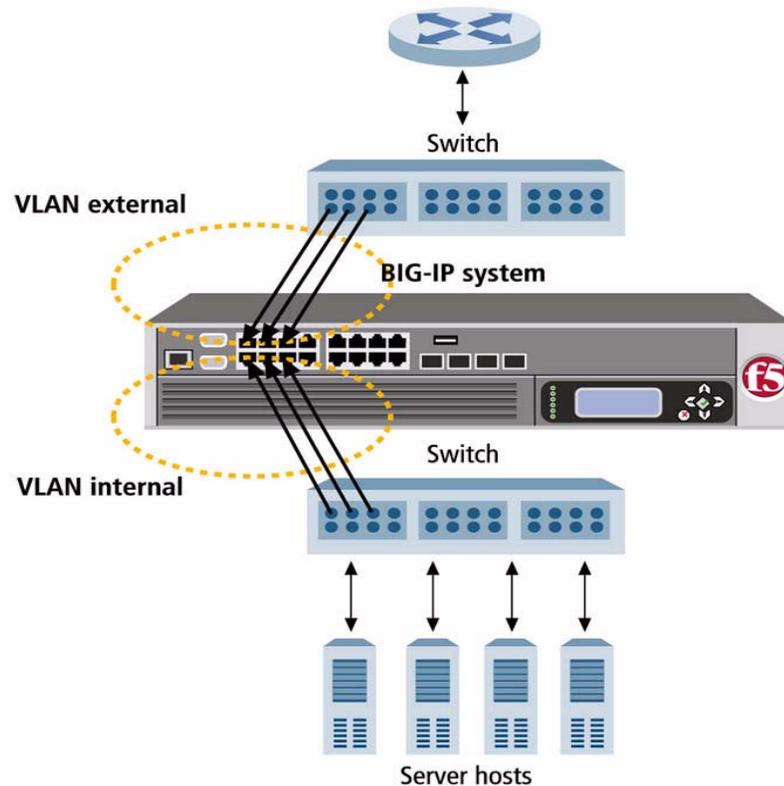


Figure 7.1 A typical configuration using the default VLANs

Every VLAN must have a static self IP address associated with it. The *self IP address* of a VLAN represents an address space, that is, the range of IP addresses pertaining to the hosts in that VLAN. When you ran the Setup utility earlier, you assigned one static self IP address to the VLAN **external**, and one static self IP address to the VLAN **internal**. When sending a

request to a destination server, the BIG-IP system can use these self IP addresses to determine the specific VLAN that contains the destination server.

For example, suppose the self IP address of VLAN **external** is **12.1.0.100**, and the self IP address of the VLAN **internal** is **11.1.0.100**, and both self IP addresses have a netmask of **255.255.0.0**. If the IP address of the destination server is **11.1.0.20**, then the BIG-IP system can compare the self IP addresses to the host's IP address to determine that the destination server is in the VLAN **internal**. This process, combined with checking the ARP cache and a VLAN's L2 forwarding table, ensures that the BIG-IP system successfully sends the request to the destination server.

Creating and managing VLANs

When you create a VLAN, you assign a name and an identifying tag to the VLAN. Then you associate one or more BIG-IP system interfaces to the VLAN. Also, if the BIG-IP system is a unit of a redundant system, you can specify a special MAC address that the two units share, as a way to ensure that connections are successfully processed when failover occurs. Finally, you can specify that you want the BIG-IP system to use VLAN-related events to trigger failover in a redundant-system configuration.

To create a VLAN, you use the Configuration utility. For information on managing an existing VLAN, see *Managing a VLAN*, on page 7-10.

Creating a VLAN

The BIG-IP system offers several settings that you can configure for a VLAN. These settings are summarized in Table 7.1.

Setting	Description	Default Value
Name	Specifies a unique name for the VLAN. This value is required.	No default value
Tag	Specifies the VLAN ID. If you do not specify a VLAN ID, the BIG-IP system assigns an ID automatically. The value of a VLAN tag can be between 1 and 4094 .	No default value
Interfaces	Specifies any tagged or untagged interfaces or trunks that you want to associate with the VLAN.	No default value
Source Check	Causes the BIG-IP system to verify that the return path of an initial packet is through the same VLAN from which the packet originated.	Unchecked
MTU	Specifies the maximum transmission unit for the VLAN.	1500

Table 7.1 Configuration settings for a VLAN

Setting	Description	Default Value
MAC Masquerade	Sets up a media access control (MAC) address that is shared by a redundant system.	No default value
Fail-safe	Triggers fail-over in a redundant system when certain VLAN-related events occur.	Unchecked

Table 7.1 Configuration settings for a VLAN

Use the following procedure to create a VLAN. For detailed information about each setting, see the sections following the procedure.

◆ Important

*In addition to configuring the settings listed in Table 7.1, you must also assign a self IP address to the VLAN. For more information, see **Assigning self IPs to VLANs and VLAN groups**, on page 7-19, and Chapter 6, **Configuring Self IP Addresses**.*

To create a VLAN

1. On the Main tab of the navigation pane, expand **Network** and click **VLANs**.
This displays a list of all existing VLANs.
2. In the upper-right corner, click **Create**.
The VLANs screen opens.
3. Locate the General Properties area, and in the **Name** box, type a unique name for the VLAN.
4. In the **Tag** box, type a tag for the VLAN, or leave the box blank.
If you do not specify a tag, the BIG-IP system assigns one automatically.
5. In the Resources area, for the **Interfaces** setting, click an interface number or trunk name in the **Available** box, and using a Move button (<< or >>), move the interface number to the **Untagged** or **Tagged** box. Repeat this step as necessary.
For more information on tagged and untagged interfaces, see *Assigning interfaces to a VLAN*, on page 7-5.
6. If you want to enable source checking, then in the Configuration area, click the **Source Check** box.
7. For the **MTU** setting, use the default value or type a new value.
8. In the **MAC Masquerade** box, type a MAC address.
For more information, see *Specifying a MAC masquerade address*, on page 7-8.

9. For the **Fail-safe** setting, check the box if you want to base redundant-system failover on VLAN-related events.
For more information, see Chapter 12, *Setting up a Redundant System*.
10. Click **Finished**.

Specifying a VLAN name

When creating a VLAN, you must assign it a unique name. Once you have finished creating the VLAN, the VLAN name appears in the list of existing VLANs.

Specifying a VLAN tag

A VLAN *tag* is a unique ID number that you assign to a VLAN. If you do not explicitly assign a tag to a VLAN, the BIG-IP system assigns a tag automatically. The value of a VLAN tag can be between **1** and **4094**. Once you or the BIG-IP assigns a tag to a VLAN, any message sent from a host in that VLAN includes this VLAN tag as a header in the message.

A VLAN tag is useful when an interface has multiple VLANs associated with it; that is, when the interfaces you assigned to the VLAN are assigned as tagged interfaces. In this case, the BIG-IP system can read the VLAN tag in the header of a message to determine the specific VLAN in which the source or destination host resides. For more information on tagged interfaces, see *Tag-based access to VLANs*, on page 7-6.

◆ Important

If the device connected to a BIG-IP system interface is another switch, the VLAN tag that you assign to the VLAN on the BIG-IP system interface must match the VLAN tag assigned to the VLAN on the interface of the other switch.

Assigning interfaces to a VLAN

For each VLAN that you create, you must assign one or more BIG-IP system interfaces to that VLAN, using the **Interfaces** setting. When you assign an interface to a VLAN, you indirectly control the hosts from which the BIG-IP system interface sends or receives messages.

◆ Tip

*You can use the **Interfaces** setting to assign not only individual interfaces to the VLAN, but also trunks. Any trunks that you create are automatically included for selection in the list of available interfaces. For more information on trunks, see Chapter 8, **Working with Trunks**.*

For example, if you assign interface 1.11 to VLAN A, and you then associate VLAN A with a virtual server, then the virtual server sends its outgoing traffic through interface 1.11, to a destination host in VLAN A.

Similarly, when a destination host sends a message to the BIG-IP system, the host's VLAN membership determines the BIG-IP system interface that should receive the incoming traffic.

Each VLAN has a MAC address. The MAC address of a VLAN is the same MAC address of the lowest-numbered interface assigned to that VLAN.

The BIG-IP system supports two methods for sending and receiving messages through an interface that is a member of one or more VLANs. These two methods are port-based access to VLANs and tag-based access to VLANs. The method used by a VLAN is determined by the way that you add a member interface to a VLAN.

Port-based access to VLANs

With port-based access to VLANs, the BIG-IP system accepts frames for a VLAN simply because they are received on an interface that is a member of that VLAN. With this method, an interface is an untagged member of the VLAN. Frames sent out through *untagged* interfaces contain no tag in their header.

Port-based access to VLANs occurs when you add an interface to a VLAN as an untagged interface. In this case, the VLAN is the only VLAN that you can associate with that interface. This limits the interface to accepting traffic only from that VLAN, instead of from multiple VLANs. If you want to give an interface the ability to accept and receive traffic for multiple VLANs, you add the same interface to each VLAN as a tagged interface. The following section describes tagged interfaces.

Tag-based access to VLANs

With tag-based access to VLANs, the BIG-IP system accepts frames for a VLAN because the frames have tags in their headers and the tag matches the VLAN identification number for the VLAN. An interface that accepts frames containing VLAN tags is a *tagged member* of the VLAN. Frames sent out through tagged interfaces contain a tag in their header.

Tag-based access to VLANs occurs when you add an interface to a VLAN as a tagged interface. You can add the same tagged interface to multiple VLANs, thereby allowing the interface to accept traffic from each VLAN with which the interface is associated.

When you add an interface to a VLAN as a tagged interface, the BIG-IP system associates the interface with the VLAN identification number, or *tag*, which becomes embedded in a header of a frame.

◆ **Note**

*Every VLAN has a tag. You can assign the tag explicitly when creating the VLAN, or the BIG-IP system assigns it automatically if you do not supply one. For more information on VLAN tags, see **Specifying a VLAN tag**, on page 7-5.*

Each time you add an interface to a VLAN, either when creating a VLAN or modifying its properties, you can designate that interface as a tagged interface. A single interface can therefore have multiple tags associated with it.

The result is that whenever a frame comes into that interface, the interface reads the tag that is embedded in a header of the frame. If the tag in the frame matches any of the tags associated with the interface, the interface accepts the frame. If the tag in the frame does *not* match any of the tags associated with the interface, the interface rejects the frame.

Example

Figure 7.2 shows the difference between using three untagged interfaces (where each interface must belong to a separate VLAN) versus one tagged interface (which belongs to multiple VLANs).

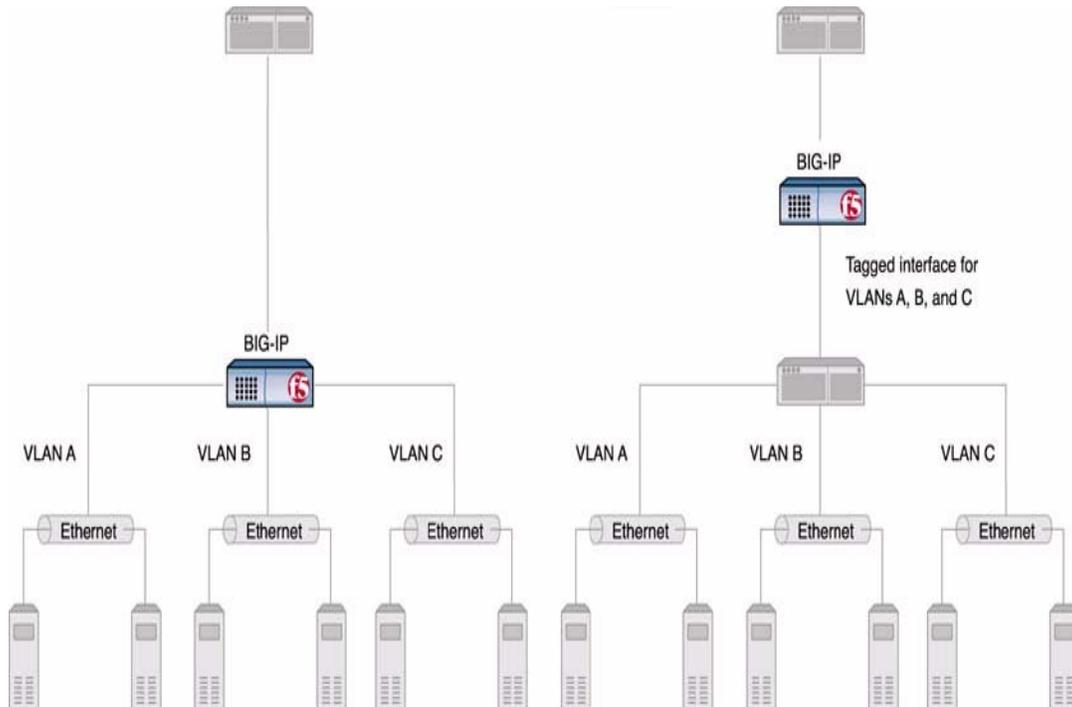


Figure 7.2 Equivalent solutions using untagged and tagged interfaces

The configuration on the left shows a BIG-IP unit with three internal interfaces, each a separate, untagged interface. This is a typical solution for supporting three separate customer sites. In this scenario, each interface can accept traffic only from its own VLAN.

Conversely, the configuration on the right shows a BIG-IP system with one internal interface and an external switch. The switch places the internal interface on three separate VLANs. The interface is configured on each VLAN as a tagged interface. In this way, the single interface becomes a tagged member of all three VLANs, and accepts traffic from all three. The configuration on the right is the functional equivalent of the configuration on the left.

◆ **Important**

If you are connecting another switch into a BIG-IP system interface, the VLAN tag that you assign to the VLAN on the BIG-IP system must match the VLAN tag on the interface of the other switch.

Enabling source checking

When you enable the **Source Check** setting, the BIG-IP system verifies that the return path for an initial packet is through the same VLAN from which the packet originated. The system performs this verification only if you check the **Source Check** box for the VLAN, and if the global setting **Auto Last Hop** is not enabled. For information on the **Auto Last Hop** setting, see Chapter 4, *Configuring the BIG-IP Platform and General Properties*.

Specifying the maximum transmission units

The value of the maximum transmission unit, or *MTU*, is the largest size that the BIG-IP system allows for an IP datagram passing through a BIG-IP system interface. The default value is **1500**.

Specifying a MAC masquerade address

Every VLAN has a media access control (MAC) address that corresponds to the VLAN's self IP address. The MAC address of a VLAN is the MAC address of the lowest-numbered interface assigned to that VLAN. For example, if the lowest-numbered interface assigned to VLAN **internal** is 3.1, and the MAC address of that interface is **0:0:0:ac:4c:a2**, then the MAC address of VLAN **internal** is also **0:0:0:ac:4c:a2**.

A **MAC masquerade address** is a variation of the VLAN's MAC address, and this address is shared between two units of a redundant system. When you specify a MAC masquerade address, a destination server sending a response to the BIG-IP system sends its response to the VLAN's MAC masquerade address, instead of to the VLAN's regular MAC address. The server accomplishes this by using the VLAN's floating self IP address as the default route when sending responses to the BIG-IP system. (For more information on configuring a server to use a floating IP address as the default route, see Chapter 12, *Setting up a Redundant System*.)

Specifying a MAC masquerade address for a VLAN has the following advantages:

- Increased reliability and failover speed, especially in lossy networks
- Interoperability with switches that are slow to respond to the network changes
- Interoperability with switches that are configured to ignore network changes

When you assign a MAC masquerade address to a VLAN, the BIG-IP system automatically sends a gratuitous ARP message to the default router and other devices on the network. This gratuitous ARP message notifies these devices that the MAC address of the BIG-IP system interface assigned to the VLAN has changed to the MAC masquerade address.

The MAC masquerade address must be a unique address, in order to avoid frame collisions. The safest way to create a MAC masquerade address is to first determine the MAC address of the VLAN (that is, the MAC address of the lowest-numbered interface assigned to that VLAN), and then logically **OR** the first byte with **0x02**. This makes the MAC address a locally-administered MAC address.

Continuing with the example above where the VLAN's MAC address is **0:0:0:ac:4c:a2**, a MAC masquerade address of **02:0:0:ac:4c:a2** is suitable to use on both BIG-IP units in the redundant system. For help in finding the MAC address of a VLAN, see *To find the MAC address of a VLAN*, on page 7-9.

Important

We highly recommend that you set the MAC masquerade address to be the same on both the active and standby units.

To find the MAC address of a VLAN

1. On the Main tab of the navigation pane, expand **Network** and click **VLANs**.
This displays a list of existing VLANs.
2. Click a VLAN name.
This displays the properties for that VLAN.
3. In the **Interfaces** setting, note the lowest-numbered interface assigned to the VLAN.
4. On the Main tab of the navigation pane, expand **Network** and click **Interfaces**.
This displays a list of all BIG-IP system interfaces and their MAC addresses.
5. Locate the interface number that you noted on the VLAN's properties screen.
6. In the MAC Address column, view the MAC address for the interface.

After you have found the correct MAC address, create the MAC masquerade address using the procedure described in step 8 in *Creating a VLAN*, on page 7-3.

Specifying fail-safe

VLAN fail-safe is a feature you enable when you want to base redundant-system failover on VLAN-related events. For more information, see Chapter 12, *Setting up a Redundant System*.

Managing a VLAN

After you have created a VLAN, you can use the Configuration utility to modify its properties, delete the VLAN, or to maintain its layer 2 forwarding table.

Managing VLAN properties

Using the Configuration utility, you can modify all of the properties of a VLAN, except the VLAN name, the tag, and the MAC address with which the VLAN is associated (that is, the MAC address of the lowest-numbered interface that is assigned to the VLAN).

You can also use the Configuration utility to delete a VLAN.

To change VLAN properties

1. On the Main tab of the navigation pane, expand **Network** and click **VLANs**.
This displays a list of all existing VLANs.
2. Click the name of the VLAN you want to modify.
This opens the properties screen for the VLAN.
3. Modify the values of any settings.
4. Click **Update**.

To delete a VLAN

1. On the Main tab of the navigation pane, expand **Network** and click **VLANs**.
This displays a list of all existing VLANs.
2. Check the Select box next to the VLAN name.
3. Click **Delete**.
A confirmation query appears.
4. Click **Delete**.

Maintaining the L2 forwarding table

Layer 2 forwarding is the means by which frames are exchanged directly between hosts, with no IP routing required. This is accomplished using a simple forwarding table for each VLAN. The **L2 forwarding table** is a list that shows, for each host in the VLAN, the MAC address of the host, along with the interface that the BIG-IP system needs for sending frames to that host. The intent of the L2 forwarding table is to help the BIG-IP system determine the correct interface for sending frames, when the system determines that no routing is required.

The format of an entry in the L2 forwarding table is:

```
<MAC address> -> <if>
```

For example, an entry for a host in the VLAN might look like this:

```
00:a0:c9:9e:1e:2f -> 2.1
```

The BIG-IP system learns the interfaces that correspond to various MAC entries as frames pass through the system, and automatically adds entries to the table accordingly. These entries are known as **dynamic entries**. You can also add entries to the table manually, and these are known as **static entries**. Entering static entries is useful if you have network devices that do not advertise their MAC addresses. The system does not automatically update static entries.

The BIG-IP system does not always need to use the L2 forwarding table to find an interface for frame transmission. For instance, if a VLAN has only one interface assigned to it, then the BIG-IP system automatically uses that interface.

Occasionally, the L2 forwarding table does not include an entry for the destination MAC address and its corresponding BIG-IP system interface. In this case, the BIG-IP system floods the frame through all interfaces associated with the VLAN, until a reply creates an entry in the L2 forwarding table.

Viewing the L2 forwarding table

You can use the Configuration utility to view the entries in the L2 forwarding table.

To view the L2 forwarding table

1. On the Main tab of the navigation pane, expand **Network** and click **VLANs**.
This displays a list of all existing VLANs.
2. Click the name of the VLAN you want to modify.
This opens the properties screen for the VLAN.
3. On the menu bar, click **Layer 2 Static Forwarding Table**.
This displays any entries currently in the L2 forwarding table.

Adding entries to the L2 forwarding table

You can add static entries to the L2 forwarding table when you want to give the BIG-IP system the ability to send messages to a specific host in the VLAN.

To add an entry to the L2 forwarding table

1. On the Main tab of the navigation pane, expand **Network** and click **VLANs**.
This displays a list of all existing VLANs.
2. Click the name of the VLAN you want to modify.
This opens the properties screen for the VLAN.
3. On the menu bar, click **Layer 2 Static Forwarding Table**.
This displays any entries currently in the L2 forwarding table.
4. Click **Create**.
This displays the screen for adding entries to the table.
5. For the **Interfaces** setting, select an interface number.
6. In the **MAC Address** box, type the MAC address of the host to which the entry applies.
7. Click **Repeat** if you want to add another entry, or click **Finished**.

Setting the L2 forwarding aging time

Entries in the L2 forwarding table have a specified life span, after which they are removed if the MAC address is no longer present on the network. This life span is called the *layer 2 cache aging time*. The default value is 300 seconds. Using the Configuration utility, you can change this value.

To change the layer 2 cache aging time

1. On the Main tab of the navigation pane, expand **System** and click **General Properties**.
This displays a list of general properties for the BIG-IP system.
2. On the menu bar, from **Local Traffic**, choose **General**.
This displays a list of general properties related to local traffic.
3. In the **Layer 2 Cache Aging Time** box, change the value.
4. Click **Update**.

Creating and managing VLAN groups

A *VLAN group* is a logical container that includes two or more distinct VLANs. VLAN groups are intended for load balancing traffic in a layer 2 network, when you want to minimize the reconfiguration of hosts on that network. Figure 7.3 shows an example of a VLAN group.

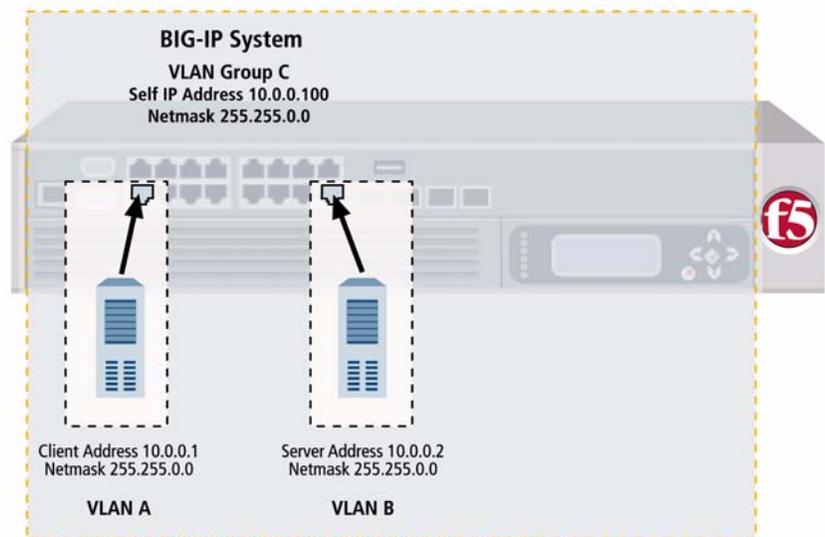


Figure 7.3 Example of a VLAN group

VLAN groups are particularly useful when you want to process traffic between a client and server that reside in the same address space, while still preserving the client IP address. (Note that *processing traffic* refers to performing load balancing or other optimizations that the BIG-IP system supports.)

Normally, if you want to process traffic between a client and server residing in the same address space, you must enable source network address translation (SNAT), which translates the client IP address to a different source address to which you want the server to respond. Although implementing a SNAT ensures that the BIG-IP system processes the traffic before returning that traffic to the client, it requires you to translate the client IP address to another source address.

A **VLAN group** also ensures that the BIG-IP system processes traffic between a client and server when the two hosts reside in the same address space. However, when you configure a VLAN group, you do not need to translate the client IP address to a different source address. You can preserve the original client IP address, and therefore no SNAT is required.

◆ **Tip**

*You can configure the behavior of the BIG-IP system so that it always creates a proxy for any ARP requests between VLANs. For more information, see **Excluding hosts from proxy ARP forwarding**, on page 7-18.*

When you create a VLAN group, the two existing VLANs become child VLANs of the VLAN group. To create a VLAN group, you use the Configuration utility. For information on managing a VLAN group, see *Managing a VLAN group*, on page 7-18.

Creating a VLAN group

When you create a VLAN group, you assign a name and a VLAN group ID. Then you specify the existing VLANs that you want the VLAN group to contain. Finally, you specify a transparency mode, and some settings related to redundant-system configuration.

◆ **Note**

Two distinct VLANs must exist on the BIG-IP system before you can create a VLAN group.

The settings that you can configure for a VLAN group are summarized in Table 7.2.

Setting	Description	Default Value
Name	Specifies a unique name for the VLAN group. This value is required.	No default value
VLAN Group ID	Specifies an ID for the VLAN group. If you do not specify a VLAN group ID, the BIG-IP system assigns an ID automatically. The value of a VLAN group ID can be between 1 and 4094 .	No default value
VLANs	Specifies the VLANs that you want the VLAN group to contain.	No default value
Transparency Mode	Specifies the level of exposure of remote MAC addresses within a VLAN group. Possible values are: Opaque , Translucent , and Transparent .	Translucent

Table 7.2 Configuration options for VLANs

Setting	Description	Default Value
Bridge All Traffic	When enabled (checked), specifies that the VLAN group forwards all frames, including non-IP traffic. The default setting is disabled (unchecked).	Enabled
Bridge in Standby	When enabled (checked), specifies that the VLAN group forwards frames, even when the system is the standby unit in a redundant system. The default setting is enabled (checked).	Disabled
MAC Masquerade	Specifies a MAC masquerade address, used when you have a redundant system.	No default value

Table 7.2 Configuration options for VLANs

Use the following procedure to create a VLAN group. For detailed information about each setting, see the sections following the procedure.

◆ Important

In addition to configuring the settings listed in Table 7.2, you must also assign a self IP address to the VLAN group. For more information, see [Assigning self IPs to VLANs and VLAN groups](#), on page 7-19.

To create a VLAN group

1. On the Main tab of the navigation pane, expand **Network** and click **VLANs**.
This displays a list of all existing VLANs.
2. On the menu bar, from **VLAN Groups**, choose List.
This displays a list of all existing VLAN groups.
3. In the upper-right corner, click **Create**.
The VLAN Groups screen opens.
4. In the General Properties area, in the **VLAN Group** box, type a unique name for the VLAN group.
5. In the **VLAN Group ID** box, type a unique VLAN ID.
If you do not specify a VLAN ID, the BIG-IP system automatically assigns one.
6. In the Configuration area, for the **VLANs** setting, click a VLAN name in the **Available** box, and using the Move button (<<), move the VLAN name to the **Members** box.
Repeat this step as necessary.
7. From the **Transparency Mode** list, select a transparency mode, or use the default setting.
8. Check the **Bridge All Traffic** setting if you want the VLAN group to forward all frames, including non-IP traffic.

9. For the **Bridge in Standby** setting, leave the box checked if you want the VLAN group to forward frames even when the system is the standby unit of a redundant system.
10. In the **MAC Masquerade** box, type a MAC address.
For more information, see *Specifying a MAC masquerade address*, on page 7-8.
11. Click **Finished**.

Specifying a VLAN group name

When creating a VLAN group, you must assign it a unique name. Once you have finished creating the VLAN group, the VLAN group name appears in the list of existing VLANs groups.

Specifying a VLAN group ID

A **VLAN group ID** is a tag for the VLAN group. Every VLAN group needs a unique ID number. If you do not specify an ID for the VLAN group, the BIG-IP system automatically assigns one. The value of a VLAN group ID can be between **1** and **4094**. For more information on VLAN tags, see *Tag-based access to VLANs*, on page 7-6.

Specifying the transparency mode

The BIG-IP system is capable of processing traffic using a combination of layer 2 and layer 3 forwarding, that is, switching and IP routing. When you set the transparency mode, you specify the type of forwarding that the BIG-IP system performs when forwarding a message to a host in a VLAN. The default setting is **translucent**, which means that the BIG-IP system uses a mix of Layer 2 and Layer 3 processing. Table 7.3 lists the allowed values.

Value	Description
opaque	A proxy ARP with layer 3 forwarding.
translucent	Layer 2 forwarding with a locally-unique bit, toggled in ARP response across VLANs. This is the default setting.
transparent	Layer 2 forwarding with the original MAC address of the remote system preserved across VLANs.

Table 7.3 Modes for VLAN group forwarding

Bridging all traffic

When you enable this option, you are instructing the VLAN group to forward all non-IP traffic. The default setting is disabled (unchecked).

Bridging traffic with standby units

When this option is enabled (checked), specifies that the VLAN group forwards frames, even when the system is the standby unit in a redundant system. The default setting is enabled (checked).

Specifying a MAC masquerade address

When you place VLANs into a VLAN group, devices on the network automatically send responses to the MAC masquerade address that you assigned to the VLAN group. In this case, the BIG-IP system ignores the MAC masquerade addresses that you assigned to the individual VLANs of the group.

The procedure for assigning a MAC masquerade address to a VLAN group is similar to the procedure for assigning one to a VLAN. However, because interfaces are not assigned directly to a VLAN group, you can use the MAC address of the lowest-numbered interface *of any VLAN in the VLAN group* when you decide on a MAC masquerade address for the VLAN group.

For more information on MAC masquerade addresses, see *Specifying a MAC masquerade address*, on page 7-8.

Managing a VLAN group

Using the Configuration utility, you can change the properties of a VLAN group, delete the VLAN group, or manage the way that the VLAN group handles proxy ARP forwarding.

Changing VLAN group properties

Using the Configuration utility, you can modify all of the properties of a VLAN group, except the VLAN name and VLAN group ID.

To change the properties of a VLAN group

1. On the Main tab of the navigation pane, expand **Network** and click **VLANs**.
This displays a list of existing VLANs.
2. From the **VLAN Groups** menu, choose List.
This shows a list of existing VLAN groups.
3. Click a VLAN group name.
4. Change the value of any settings.
5. Click **Update**.

Excluding hosts from proxy ARP forwarding

As described earlier, a host in a VLAN cannot normally communicate to a host in another VLAN. This rule applies to ARP requests as well. However, if you put the VLANs into a single VLAN group, the BIG-IP system can perform a proxied ARP request.

A *proxied ARP request* is an ARP request that the BIG-IP system can send, on behalf of a host in a VLAN, to hosts in another VLAN. A proxied ARP request requires that both VLANs belong to the same VLAN group.

In some cases, you might not want a host to forward proxied ARP requests to a specific host, such as an active unit in a redundant system that forwards a proxied ARP request to the standby unit, or to other hosts in the configuration. To exclude specific hosts from receiving forwarded proxied ARP requests, you use the Configuration utility and specify the IP addresses that you want to exclude.

WARNING

Although hosts on an ARP exclusion list are specified using their IP addresses, this does not prevent the BIG-IP system from routing traffic to those hosts. A more secure way to prevent traffic from passing between hosts in separate VLANs is to create a packet filter for each VLAN.

To exclude the forwarding of proxied ARP requests

1. On the Main tab of the navigation pane, expand **Network** and click **VLANs**.
This displays a list of all existing VLANs.
2. On the menu bar, from **VLAN Groups**, choose Proxy Exclusion List.
This opens the Global Proxy Exclusion List screen.
3. In the upper-right corner, click **Create**.
4. In the IP address box, type an IP address that you want to exclude from a proxied ARP request.
5. Click **Repeat** if you want to type another IP address, or click **Finished**.

Assigning self IPs to VLANs and VLAN groups

After you create a VLAN or a VLAN group, you must assign it a self IP address. You assign self IP addresses to VLANs and VLAN groups using the Configuration utility.

◆ Assigning a self IP address to a VLAN

The self IP address that you assign to a VLAN should represent an address space that includes the self IP addresses of the hosts that the VLAN contains. For example, if the address of one host is **11.0.0.1** and the address of the other host is **11.0.0.2**, you could assign an address of **11.0.0.100**, with a netmask of **255.255.255.0**, to the VLAN.

◆ Assigning a self IP address to the VLAN group

The self IP address that you assign to a VLAN group should represent an address space that includes the self IP addresses of the VLANs that you assigned to the group. For example, if the address of one VLAN is **10.0.0.1** and the address of the other VLAN is **10.0.0.2**, you could assign an address of **10.0.0.100**, with a netmask of **255.255.255.0**, to the VLAN group.

For more detailed information and the procedure for assigning self IP addresses, see Chapter 6, *Configuring Self IP Addresses*.



8

Working with Trunks

- Introducing trunks
- Creating a trunk
- Managing trunks

Introducing trunks

A **trunk** is a logical grouping of interfaces on the BIG-IP system. When you create a trunk, this logical group of interfaces functions as a single interface. The BIG-IP system uses a trunk to distribute traffic across multiple links, in a process known as **link aggregation**. With link aggregation, a trunk increases the bandwidth of a link by adding the bandwidth of multiple links together. For example, four fast Ethernet (100 Mbps) links, if aggregated, create a single 400 Mbps link.

With one trunk, you can aggregate a maximum of eight links. For optimal performance, you should aggregate links in powers of two. Thus, you ideally aggregate two, four, or eight links.

The purpose of a trunk is two-fold: To increase bandwidth without upgrading hardware, and to provide link failover if a member link becomes unavailable.

You can use trunks to transmit traffic from a BIG-IP system to another vendor switch. Two systems that use trunks to exchange frames are known as **peer systems**.

How do trunks work?

In a typical configuration where trunks are configured, the member links of the trunk are connected through Ethernet cables to corresponding links on a peer system. Figure 8.1 shows an example of a typical trunk configuration with two peers and three member links on each peer.

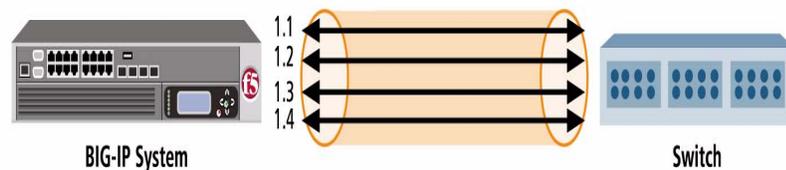


Figure 8.1 Example of a trunk configured for two switches

A primary goal of the trunks feature is to ensure that frames exchanged between peer systems are never sent out of order or duplicated on the receiving end. The BIG-IP system is able to maintain frame order by using the source and destination addresses in each frame to calculate a hash value, and then transmitting all frames with that hash value on the same member link.

The BIG-IP system automatically assigns a unique MAC address to a trunk. This MAC address becomes the source address for frames that the system transmits, and the destination address for frames that the system receives.

However, the BIG-IP system still uses the MAC address of an individual member link as the source address for any LACP control frames. For more information on LACP, see *Overview of LACP*, following.

The BIG-IP system uses the lowest-numbered interface in a trunk as a **reference link**. The BIG-IP system uses the reference link to take certain aggregation actions, such as implementing the automatic link selection policy. For frames coming into the reference link, the BIG-IP system load balances the frames across all member links that the BIG-IP system knows to be available. For frames going from any link in the trunk to a destination host, the BIG-IP system treats those frames as if they came from the reference link.

Overview of LACP

A key aspect of trunks is Link Aggregation Control Protocol, or LACP. Defined by IEEE standard 802.3ad, **LACP** is a protocol that detects error conditions on member links and redistributes traffic to other member links, thus preventing any loss of traffic on the failed link. On a BIG-IP system, LACP is an optional feature that you can configure.

You can also customize LACP behavior. For example, you can specify the way that LACP communicates its control messages from the BIG-IP system to a peer system. You can also specify the rate at which the BIG-IP system exchanges LACP packets with a peer system. If you want to affect the way that the BIG-IP system chooses links for link aggregation, you can specify a link control policy. For more information, see *Creating a trunk*, on page 8-3.

Creating a trunk

You create a trunk using the Configuration utility. The BIG-IP system offers several settings that you can configure for a trunk. These settings are summarized in Table 8.1.

Setting	Description	Default Value
Name	Specifies a unique name for the trunk. For more information, see <i>Specifying a trunk name</i> , on page 8-4.	No default value
Interfaces	Specifies the interfaces that constitute member links of the trunk. For more information, see <i>Specifying interfaces for a trunk</i> , on page 8-4.	No default value
LACP	Enables or disables the Link Aggregation Control Protocol (LACP). For more information, see <i>Enabling LACP</i> , on page 8-5.	Unchecked
Link Mode	Specifies the way that member links communicate over LACP. Possible values are Active and Passive . For more information, see <i>Specifying the LACP mode</i> , on page 8-5.	Active
Link Timeout	Specifies the length of time that the BIG-IP system sends control packets. Possible values are Short and Long . For more information, see <i>Specifying the LACP timeout</i> , on page 8-6.	Long
Link Selection Policy	Specifies the policy that the BIG-IP uses to select a link for processing traffic. Possible values are Auto and Maximum Bandwidth . For more information, see <i>Specifying a link selection policy</i> , on page 8-6.	Auto

Table 8.1 Configuration settings for a trunk

Use the following procedure to create a trunk. For detailed information about each setting, see the sections following the procedure.

To create a trunk

1. On the Main tab, expand **Network**, and click **Trunks**. The Trunks screen opens.
2. In the upper-right corner of the screen, click the **Create** button. Another Trunks screen opens.
3. In the **Name** box, type a unique name for the trunk.
4. For the **Interfaces** setting, in the **Available** box, select an interface that you want to include in the trunk, and using the Move button (<<), move the interface to the **Members** box.
5. Repeat the previous step as necessary.

6. If you want to enable the LACP feature, check the **LACP** box, and do the following:
 - a) From the **LACP Mode** list, select **Passive** mode, or retain the default mode (**Active**).
 - b) From the **LACP Timeout** list, select **Short**, or retain the default timeout value (**Long**).
7. From the **Link Selection Policy** list, select or retain a link selection policy.
8. Click **Finished**.

Specifying a trunk name

You can use the **Name** setting to specify a unique name for the trunk. This setting is required.

Specifying interfaces for a trunk

Using the **Interfaces** setting, you specify the interfaces that you want the BIG-IP system to use as member links for the trunk. Once you have created the trunk, the BIG-IP system uses these interfaces to perform link aggregation.

As stated earlier, the BIG-IP system uses the lowest-numbered interface as the reference link. The system uses the reference link to negotiate links for aggregation. For more information on link negotiation, see *Specifying a link selection policy*, on page 8-6.

The interfaces that you specify for the trunk must operate at the same media speed, and must be set at full-duplex mode. Otherwise, the BIG-IP system cannot aggregate the links. Because these media properties are dynamic rather than static (due to auto-negotiation), the **lacpd** service routinely monitors the current status of these properties and negotiates the links for aggregation accordingly. Thus, when the status of these properties qualifies a link to become a working member link, the system adds the link to the aggregation, and the link can begin accepting traffic. For information on setting media properties for an interface, see *Platform Guide: 1500, 3400, 6400, and 6800*.

Any interface that you assign to a trunk must be an untagged interface. Furthermore, you can assign an interface to one trunk only; that is, you cannot assign the same interface to multiple trunks. Because of these restrictions, the only interfaces that appear in the **Interfaces** list in the Configuration utility are untagged interfaces that are not assigned to another trunk. Therefore, before creating a trunk and assigning any interfaces to it, you should verify that each interface for the trunk is an untagged interface.

After creating the trunk, you assign the trunk to one or more VLANs, using the same VLAN screen that you normally use to assign an individual interface to a VLAN. For information on assigning a trunk to a VLAN, see Chapter 7, *Configuring VLANs and VLAN Groups*.

If you are using one of the spanning tree protocols (STP, RSTP, or MSTP), the BIG-IP system sends and receives spanning tree protocol packets on a trunk, rather than on individual member links. Likewise, use of a spanning tree protocol to enable or disable learning or forwarding on a trunk operates on all member links together, as a single unit.

Enabling LACP

As an option, you can enable LACP on a trunk. Containing a service called **lacpd**, LACP is an IEEE-defined protocol that exchanges control packets over member links. The purpose of LACP is to detect link error conditions such as faulty MAC devices and link loopbacks. If LACP detects an error on a member link, the BIG-IP system removes the member link from the link aggregation and redistributes the traffic for that link to the remaining links of the trunk. In this way, no traffic destined for the removed link is lost. LACP then continues to monitor the member links to ensure that aggregation of those links remains valid.

By default, the LACP feature is disabled, to ensure backward compatibility with previous versions of the BIG-IP system. If you create a trunk and do not enable the LACP feature, the BIG-IP system does not detect link error conditions, and therefore cannot remove the member link from link aggregation. The result is that the system cannot redistribute the traffic destined for that link to the remaining links in the trunk, thereby causing traffic on the failed member link to be lost.

◆ Important

To use LACP successfully, you must enable LACP on both peer systems.

Specifying the LACP mode

The **LACP Mode** setting appears on the Trunks screen only when you check the **LACP** setting. You use the **LACP mode** setting to specify the method that LACP uses to send control packets to the peer system. The two possible modes are:

◆ Active mode

You specify **Active** mode if you want the system to periodically send control packets, regardless of whether the peer system has issued a request. This is the default setting.

◆ Passive mode

You specify **Passive** mode if you want the system to send control packets only when the peer system issues a request, that is, when the LACP mode of the peer system is set to **Active**.

If you set only one of the peer systems to **Active** mode, the BIG-IP system uses **Active** mode for both systems. Also, whenever you change the LACP mode on a trunk, LACP renegotiates the links that it uses for aggregation on that trunk.

◆ **Tip**

*We recommend that you set the LACP mode to **Passive** on one peer system only. If you set both systems to **Passive** mode, no control packets are sent.*

Specifying the LACP timeout

The **LACP Timeout** setting appears on the Trunks screen only when you check the **LACP** setting. You use the **LACP Timeout** setting to specify the number of seconds that the system waits before sending control packets to the peer system. The timeout value only applies when the LACP mode is set to **Active** on at least one of the systems. If both systems are set to **Passive** mode, LACP does not send control packets.

If LACP sends three consecutive control packets without receiving a response from the peer system, LACP removes that member link from link aggregation.

The two possible timeout values are:

◆ **Short**

When you set the timeout value to **Short**, the BIG-IP system sends LACP control packets to the peer system once every second. If this value is set to **Short** and LACP receives no peer response after sending three consecutive packets, LACP removes the link from aggregation in three seconds.

◆ **Long**

When you set the timeout value to **Long**, the BIG-IP system sends LACP control packets to the peer system once every 30 seconds. A timeout value of **Long** is the default setting. If set to **Long** and LACP receives no peer response after sending three consecutive packets, LACP removes the link from aggregation in ninety seconds.

Whenever you change the LACP timeout value on a trunk, LACP renegotiates the links that it uses for aggregation on that trunk.

Specifying a link selection policy

In order for the BIG-IP system to aggregate links, the media speed and duplex mode of each link must be the same on both peer systems. Because media properties can change dynamically, the BIG-IP system monitors these properties regularly, and if it finds that the media properties of a link are mismatched on the peer systems, the BIG-IP system must determine which links are eligible for aggregation.

The way the system determines eligible links depends on a link selection policy that you choose for the trunk. When you create a trunk, you can choose one of two possible policy settings: **Auto** and **Maximum Bandwidth**.

◆ **Note**

The link selection policy feature represents an F5 Networks enhancement to the standard IEEE 802.3ad specification for LACP.

Understanding automatic link selection

When you set the link selection policy to **Auto** (the default setting), the BIG-IP system uses the lowest-numbered interface of the trunk as a reference link. (A **reference link** is a link that the BIG-IP uses to make a link aggregation decision.) The system then aggregates any links that have the same media properties and are connected to the same peer as the reference link.

For example, using Figure 8.1, on page 8-1, suppose that you created a trunk to include interfaces 1.2 and 1.3, each with a media speeds of 100 Mbps, and interface 1.4, with a different media speed of 1 Gbps. If you set the link selection policy to **Auto**, the BIG-IP system uses the lowest-numbered interface, 1.2, as a reference link. The reference link operates at a media speed of 100 Mbps, which means that the system aggregates all links with that media speed (interfaces 1.2 and 1.3). The media speed of interface 1.4 is different (1 Gbps), and therefore is not considered for link aggregation. Only interfaces 1.2 and 1.3 become working member links and start carrying traffic.

If the media speed of interface 1.4 changes to 100 Mbps, the system adds that interface to the aggregation. Conversely, if the media speed of interface 1.4 remains at 1 Gbps, and the speed of the reference link changes to 1 Gbps, then interfaces 1.2 and 1.4 become working members, and 1.3 is now excluded from the aggregation and no longer carries traffic.

Understanding maximum bandwidth link selection

When you set the link selection policy to **Maximum Bandwidth**, the BIG-IP system aggregates the subset of member links that provide the maximum amount of bandwidth to the trunk.

Continuing with our previous example, if interfaces 1.2 and 1.3 each operate at a media speed of 100 Mbps, and interface 1.4 operates at speed of 1 Gbps, then the system selects only interface 1.4 as a working member link, providing 1 Gbps of bandwidth to the trunk. If the speed of interface 1.4

drops to 10 Mbps, the system then aggregates links 1.2 and 1.3, to provide a total bandwidth to the trunk of 200 Mbps. The peer system detects any non-working member links and configures its aggregation accordingly.

◆ **Tip**

To ensure that link aggregation operates properly, make sure that both peer systems agree on the link membership of their trunks.

Managing trunks

After you have created your trunks, you can manage them in several ways. You can:

- View a list of existing trunks
- View or modify trunk properties
- Add a trunk to a VLAN
- Delete a trunk
- Manage interfaces for a trunk

As with trunk creation, you manage trunks using the Configuration utility.

Viewing a list of trunks

Before you create a new trunk, you might want to view a list of existing trunks. When you use the Configuration utility to view a list of trunks, you not only see the trunk names, but also some data about each trunk:

- Trunk status (**UP** or **DOWN**)
- MAC address of the trunk
- LACP mode (**Active** or **Passive**)
- Number of interfaces (member links) in the trunk
- Current trunk capacity

To view a list of trunk names and data

1. On the Main tab, expand **Network**, and click **Trunks**. The Trunks screen opens.
2. View the list of trunk names and associated data.

Viewing or modifying trunk properties

For any existing trunk, you can use the Configuration utility to view or modify the properties of that trunk. These properties consist of not only the settings that you configured for the trunk, but also other read-only LACP-specific properties that conform to the IEEE 802.3ad standard.

The read-only LACP properties you can view for an existing trunk are as follows. Note that the term *actor* refers to the local endpoint of the trunk, and *peer* refers to the remote endpoint of the trunk.

- MAC address of the trunk
- Actor ID, Actor Key, and Actor Priority
- Peer ID, Peer Key, and Peer Priority

For detailed information these LACP parameters, see the IEEE 802.3ad specification.

The modifiable properties that you can view for a trunk are:

- Interfaces that are member links of the trunk
- LACP settings (such as mode and timeout)
- Link selection policy

To view or modify properties of a trunk

1. On the Main tab, expand **Network**, and click **Trunks**.
The Trunks screen opens, showing the list of existing trunks.
2. Click a trunk name.
The properties screen for that trunk opens.
3. In the Properties area of the screen, view the properties.
4. In the Configuration area of the screen, view or modify the setting values.
For information on selecting values for these settings, see *Creating a trunk*, on page 8-3.
5. If you modified any values, click **Update**. Otherwise, click **Cancel**.

Adding a trunk to a VLAN

After creating a trunk, you can add the trunk to a VLAN. You add a trunk to a VLAN by using the VLAN screens of the Configuration utility. For more information, see Chapter 7, *Configuring VLANs and VLAN Groups*.

Deleting a trunk

You can use the Configuration utility to delete an existing trunk. When you delete a trunk, individual interfaces are no longer considered to be member links of a trunk and are no longer aggregated.

Remember that when you added the interfaces to the trunk, the BIG-IP system first required you to remove the individual interfaces from VLAN membership. Therefore, after deleting the trunk, if you want the individual interfaces of the trunk to become VLAN members again, you must explicitly assign them to one or more VLANs, using the VLANs screens of the Configuration utility.

You cannot delete any trunk that is a member of VLAN or that participates in Spanning Tree Protocol. For this reason, you must remove the trunk from any VLAN to which it is assigned, and from STP configuration, before you delete the trunk.

To delete a trunk

1. On the Main tab, expand **Network**, and click **Trunks**.
The Trunks screen opens, showing the list of existing trunks.
2. In the Name column, locate the name of the trunk you want to delete.
3. To the left of the trunk name, check the Select box.
4. Click the **Delete** button.
A confirmation message appears.
5. Click **Delete** to delete the trunk.

Managing interfaces for a trunk

You can manage the interfaces of a trunk by adding them to or deleting them from a trunk. You can also view trunk-related properties of an interface in the trunk.

◆ Note

*The following procedure for adding an interface to a trunk is equivalent to modifying the **Interfaces** property of the trunk, as described in step 4 of the section **To view or modify properties of a trunk**, on page 8-10.*

To add an interface to a trunk by viewing trunk resources

1. On the Main tab, expand **Network**, and click **Trunks**.
The Trunks screen opens, showing the list of existing trunks.
2. Click a trunk name.
The properties screen for that trunk opens.
3. On the menu bar, click **Resources**.
This displays the list of interfaces assigned to the trunk, along with some interface properties.
4. In the upper right corner of the screen, click **Add**.
5. For the **Interfaces** setting, in the **Available** box, select an interface that you want to include in the trunk, and using the Move button (<<), move the interface to the **Members** box.
6. Repeat the previous step as necessary.
7. Click **Update**.

To remove an interface from a trunk

1. On the Main tab, expand **Network**, and click **Trunks**.
The Trunks screen opens, showing the list of existing trunks.
2. Click a trunk name.
The properties screen for that trunk opens.

3. On the menu bar, click **Resources**.
This displays the list of interfaces assigned to the trunk, along with some interface properties.
4. In the Name column, locate the interface that you want to remove, and check the Select box to the left of the Status column.
5. Click **Delete**.
A confirmation message appears.
6. Click **Delete** to remove the interface from the trunk.

To view trunk-related properties of an interface

1. On the Main tab, expand **Network**, and click **Trunks**.
The Trunks screen opens, showing the list of existing trunks.
2. Click a trunk name.
The properties screen for that trunk opens.
3. On the menu bar, click **Resources**.
This displays the list of interfaces assigned to the trunk, along with some interface properties.
4. In the Name column, click the interface for which you want to view trunk-related properties.
This displays trunk-related properties for the interface.

◆ Tip

*To view the set of properties for an interface, see Chapter 5, **Working with Interfaces**.*



9

Configuring Spanning Tree Protocols

- Introducing spanning tree protocols
- Configuring global spanning tree properties
- Managing spanning tree instances
- Configuring interfaces for spanning tree

Introducing spanning tree protocols

On networks that contain redundant paths between layer 2 devices, a common problem is bridging loops. Bridging loops occur because layer 2 devices do not create boundaries for broadcasts or packet floods. Consequently, layer 2 devices can use redundant paths to forward the same frames to each other continuously, eventually causing the network to fail.

To solve this problem, the BIG-IP system supports a set of industry-standard, layer 2 protocols known as spanning tree protocols. *Spanning tree protocols* block redundant paths on a network, thus preventing bridging loops. If a blocked, redundant path is needed later because another path has failed, the spanning tree protocols clear the path again for traffic. The spanning tree protocols that the BIG-IP system supports are Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP).

Central to the way that spanning tree protocols operate is the use of bridge protocol data units (BPDUs). When you enable spanning tree protocols on layer 2 devices on a network, the devices send *BPDUs* to each other, for the purpose of learning the redundant paths and updating their L2 forwarding tables accordingly, electing a root bridge, building a spanning tree, and notifying each other about changes in interface status.

◆ Note

*Throughout this chapter, the term **bridge** refers to a layer 2 device such as a switch, bridge, or hub.*

Spanning tree protocol types

The BIG-IP system supports three different spanning tree protocols: STP, RSTP, and MSTP. Table 9.1 lists the protocols and their IEEE specifications. Following the table is a brief summary of each protocol.

Protocol Name	IEEE Specification
Spanning Tree Protocol (STP)	802.1D-1998
Rapid Spanning Tree Protocol (RSTP)	802.1w, 802.1t, and 802.1D-2004
Multiple Spanning Tree Protocol (MSTP)	802.1s

Table 9.1 *The spanning tree protocols that the BIG-IP system supports*

The STP protocol

STP is the original spanning tree protocol, designed to block redundant paths as a way to prevent bridging loops. The STP algorithm creates one, and only one, spanning tree for the entire network. A *spanning tree* is a logical tree-like depiction of the bridges on a network and the paths that connect them.

Because STP is unable to recognize VLANs and usually exhibits poor performance overall, STP is not the preferred spanning tree protocol to use in VLAN-rich environments. However, all participating interfaces in the spanning tree must use the same spanning tree protocol at any given time. Thus, when you have legacy bridges in your environment that are running STP, interfaces on the BIG-IP system must have the ability to automatically degrade to STP. For more information on protocol degradation, see *Using spanning tree with legacy bridges*, on page 9-3.

Because STP has no knowledge of VLANs, you can have only one spanning tree instance on the BIG-IP system when using STP. For more information on spanning tree instances, see *Managing spanning tree instances*, on page 9-10.

The RSTP protocol

RSTP is an enhancement to STP, and was designed specifically to improve spanning tree performance. Like STP, RSTP can create only one spanning tree (instance 0), and therefore cannot take VLANs into account when managing redundant paths. However, RSTP's performance improvements generally make it preferable to STP in non-VLAN environments.

In the case where legacy RSTP bridges are on the network, BIG-IP system interfaces running MSTP can degrade to RSTP, just as they can degrade to STP. For more information on protocol degradation, see *Using spanning tree with legacy bridges*, on page 9-3.

Like STP, RSTP allows only one spanning tree instance on the BIG-IP system. For more information on spanning tree instances, see *Managing spanning tree instances*, on page 9-10.

The MSTP protocol

MSTP is an enhancement to RSTP and is the preferred spanning tree protocol for the BIG-IP system. MSTP is specifically designed to understand VLANs and VLAN tagging (specified in IEEE 802.1q). Unlike STP and RSTP, which allow only one spanning tree instance per system, MSTP allows multiple spanning tree instances. Each instance corresponds to a spanning tree, and can control one or more VLANs that you specify when you create the instance. Thus, for any BIG-IP system interface that you assigned to multiple VLANs, MSTP can block a path on one VLAN, while still keeping a path in another VLAN open for traffic. Neither STP nor RSTP has this capability.

A unique feature of MSTP is the concept of spanning tree regions. A **spanning tree region** is a logical set of bridges on the network that share the same values for certain MSTP configuration settings. These configuration settings are: The MSTP configuration name, the MSTP configuration number, the instance numbers, and the VLAN members of each instance. When the values of these settings are identical on two or more bridges, the spanning tree algorithm considers these bridges to constitute an MSTP region. An **MSTP region** indicates to the spanning tree algorithm that it can use MSTP for all bridges in that region, and thus take VLANs into account when blocking and unblocking redundant paths.

You do not explicitly create a region. The spanning tree algorithm automatically groups bridges into regions, based on the values you assign to the MSTP configuration name, revision number, instance numbers, and instance members.

MSTP can only operate on bridges that are within a region. However, if the BIG-IP system connects to a bridge in a different MSTP region or outside of an MSTP region, the system still participates in spanning tree. In this case, the system is part of the spanning tree instance **0**, also known as the Common and Internal Spanning Tree (CIST).

Using spanning tree with legacy bridges

A key concept about spanning tree protocols on the BIG-IP system is the concept of protocol degradation. **Protocol degradation** occurs when the spanning tree mode on the BIG-IP system is set to MSTP or RSTP, but the system detects legacy bridges (that is, bridges running an older protocol type) on the network. In this case, the BIG-IP system automatically degrades the spanning tree protocol that is running on each applicable interface to match the protocol running on the legacy device.

For example, suppose you set the BIG-IP system to run in MSTP mode. Later, if a bridge running STP is added to the network, the BIG-IP system will detect the legacy device and automatically degrade the protocol running on the BIG-IP system interfaces from MSTP to STP. The mode is still set to MSTP, but the interfaces actually run STP.

If the legacy device is later removed from the network, you can choose, for each BIG-IP system interface, to manually reset the spanning tree protocol back to MSTP.

The basic principle of protocol degradation is that each BIG-IP system interface in a spanning tree runs the oldest protocol that the system detects on the layer 2 devices of the network. Thus, if a legacy bridge running STP is added to the network, BIG-IP system interfaces running MSTP or RSTP degrade to STP. Similarly, if a legacy bridge is running RSTP (and no bridges are running STP), interfaces running MSTP degrade to RSTP.

Note that when a bridge running MSTP must degrade to RSTP, the spanning tree algorithm automatically puts the degraded bridge into a separate MSTP region.

Configuration overview

Regardless of which spanning tree protocol you choose to use, the BIG-IP system offers a complete set of default configuration settings. Except for choosing a preferred spanning tree protocol to use, there are very few configuration settings that you need to modify to use the spanning tree feature effectively.

When you configure spanning tree on a BIG-IP system, you must first decide which protocol, or mode, you want to enable. Because MSTP recognizes VLANs, using MSTP is preferable for the BIG-IP system. However, all bridges in a network environment that want to use spanning tree must run the same spanning tree protocol. If a legacy bridge running RSTP or STP is added to the network, the BIG-IP system must switch to that same protocol.

Fortunately, you do not need to continually reconfigure the BIG-IP system spanning tree mode whenever a legacy bridge is added to the network. Instead, a BIG-IP system interface can detect the addition of a legacy bridge and automatically fall back to either RSTP or STP mode. If the legacy bridge is later removed from the network, you can use the Configuration utility to manually reset the interface back to running MSTP. For more information on legacy bridges, see *Using spanning tree with legacy bridges*, on page 9-3.

Once you have enabled a spanning tree mode, you can configure a set of global options. These options are the same options that are defined in the IEEE standards for the spanning tree protocols. While you can use the default settings in most cases, a few settings require user input. For more information, see *Configuring global spanning tree properties*, on page 9-5.

Configuring global spanning tree properties

There are several properties you can configure on the BIG-IP system that affect the behavior of all spanning tree protocols. These global properties apply to all spanning instances and all network interfaces. In most cases, you can use the default values for these properties. Table 9.2 lists these global properties.

Property	Description	Default Value
Mode	Specifies the protocol you want to use or not use. Possible settings are: Disabled , Pass Through , STP , RSTP , MSTP .	Pass Through
Hello Time	Specifies, in seconds, how often the system broadcasts HELLO frames to other members of the spanning tree.	2
Maximum Age	Specifies, in seconds, the length of time for which spanning tree information from other bridges is considered valid.	20
Forward Delay	Specifies, in seconds, the length of time for which an interface is blocked from forwarding network traffic after the spanning tree topology has been modified. This property is more useful for STP than RSTP or MSTP.	15
Transmit Hold Count	Specifies the maximum number of spanning tree frames the system can transmit on a port within the Hello Time interval.	6
MSTP Configuration Name	For MSTP only, specifies the name of the spanning tree configuration. All bridges with the same MSTP configuration name, MSTP configuration revision number, instance numbers, and instance members are considered to be in the same MSTP region.	MAC address of lowest-numbered interface
MSTP Configuration Revision	For MSTP only, specifies the revision level of an MSTP configuration.	0
MSTP Maximum Hops	For MSTP only, specifies the maximum number of hops that a spanning tree frame can traverse before it is discarded.	20

Table 9.2 Global spanning tree properties

Use the following procedure to configure global spanning tree properties. For detailed information on each property, see these sections:

- *Specifying the spanning tree mode*, on page 9-6
- *Configuring global timers*, on page 9-7
- *Specifying the Transmit Hold Count option*, on page 9-8
- *Configuring MSTP-specific global properties*, on page 9-8

To configure global spanning tree properties

1. On the Main tab of the navigation pane, expand **Network**, and click **Spanning Tree**.
This opens the Spanning Tree screen.
2. On the menu bar, click **Options**.
This displays the screen for configuring global spanning tree properties.
3. Configure the properties as needed.
For information on each property, see the following sections.
4. Click **Update**.

Specifying the spanning tree mode

The **Mode** option specifies the particular spanning tree protocol that you want to use on the BIG-IP system. The default value is **Pass Through**. The possible values are:

- ◆ **Disabled**
Specifies that when the BIG-IP system receives spanning tree frames (BPDUs), it discards the frames.
- ◆ **Pass Through**
Specifies that when the BIG-IP system receives spanning tree frames (BPDUs), it forwards them to all other interfaces. This is the default setting. When you use **Pass Through** mode, the BIG-IP system is transparent to spanning tree BPDUs. When set to **Pass Through** mode, the BIG-IP system is not part of any spanning tree. Note that **Pass Through** mode is not part of the IEEE spanning tree protocol specifications.
- ◆ **STP**
Specifies that the BIG-IP system handles spanning tree frames (BPDUs) in accordance with the STP protocol. This mode allows for legacy systems on the network. For more information on STP, see *Introducing spanning tree protocols*, on page 9-1.
- ◆ **RSTP**
Specifies that the BIG-IP system handles spanning tree frames (BPDUs) in accordance with the RSTP protocol. For more information RSTP, see *Introducing spanning tree protocols*, on page 9-1.
- ◆ **MSTP**
Specifies that the BIG-IP system handles spanning tree frames (BPDUs) in accordance with the MSTP protocol. For more information MSTP, see *Introducing spanning tree protocols*, on page 9-1.

◆ Important

If you select MSTP mode, additional options appear on the screen.

When you set the mode to MSTP or RSTP, and a legacy bridge running STP is subsequently added to the spanning tree, the applicable BIG-IP system interface automatically changes to running STP. However, you can manually reset an interface to resume operation in RSTP or MSTP mode if the legacy bridge is later removed from the spanning tree. For information on detecting the protocol version, see *Configuring interfaces for spanning tree*, on page 9-16.

Configuring global timers

All three spanning tree protocols, have the same three global timer values that you can specify: **Hello Time**, **Maximum Age**, and **Forward Delay**.

Specifying the Hello Time option

When you change the value of the **Hello Time** option, you change the time interval, in seconds, that the BIG-IP system transmits spanning tree information (through BPDUs) to adjacent bridges in the network. The default value for this option is **2**.

WARNING

Although valid values are in the range of 1 to 10 seconds, we highly recommend that you use the default value (2 seconds). This value is optimal for almost all configurations.

Note that when running RSTP, you must maintain the following relationship between the **Maximum Age** and **Hello Time** options:

```
Maximum Age >= 2 * (Hello Time + 1)
```

Specifying the Maximum Age option

When you change the value of the **Maximum Age** option, you change the amount of time, in seconds, that spanning tree information received from other bridges is considered valid. The default value is **20**, and the valid range is 6 to 40.

Note that when running RSTP, you must maintain the following relationships between the **Maximum Age** and the **Hello Time** and **Forward Delay** options:

```
Maximum Age >= 2 * (Hello Time + 1)
```

```
Maximum Age <= 2 * (Forward Delay - 1)
```

Specifying the Forward Delay option

Primarily used for STP, the **Forward Delay** option specifies the amount of time, in seconds, that the system blocks an interface from forwarding network traffic when the spanning tree algorithm reconfigures a spanning tree. The default value is **15**, and the valid range is 4 to 30.

This option has no effect on the BIG-IP system when running RSTP or MSTP mode, as long as all bridges in the spanning tree use the RSTP or MSTP protocol. However, if the addition of legacy STP bridges causes neighboring bridges to fall back to running the STP protocol, then the spanning tree algorithm uses the **Forward Delay** option when reconfiguring the spanning tree.

Note that when running RSTP, you must maintain the following relationship between the **Forward Delay** and **Maximum Age** options:

Maximum Age $\leq 2 * (\text{Forward Delay} - 1)$

Specifying the Transmit Hold Count option

When you change the value of the **Transmit Hold Count** option, you change the maximum number of spanning tree frames (BPDUs) that the system can transmit on a port within the **Hello Time** interval. This setting ensures that the spanning tree frames do not overload the network, even in unstable network conditions. The default value is **6**, and the valid range is 1 to 10.

Configuring MSTP-specific global properties

If you are running MSTP, you can configure three additional global properties: An MSTP configuration name, an MSTP configuration revision, and a maximum hop number.

Specifying an MSTP configuration name

Applicable to MSTP only, the **MSTP Configuration Name** setting represents a global name that you assign to all bridges in a spanning tree region. A *spanning tree region* is a group of bridges with identical MSTP configuration names and MSTP configuration revision levels, as well as identical assignment of VLANs to spanning tree instances.

All bridges in the same region must have this same configuration name. The name must contain from 1 to 32 characters. This option only appears on the screen when you set the **Mode** property to **MSTP**.

For more information on MSTP regions, see *The MSTP protocol*, on page 9-2.

Specifying an MSTP configuration revision

Applicable to MSTP only, the **MSTP Configuration Revision** setting represents a global revision number that you assign to all bridges in a spanning tree region. All bridges in the same region must have this same configuration revision number. The default value is **0**. You can type any value between 0 and 65535. This option only appears on the screen when you set the **Mode** property to **MSTP**.

For more information on MSTP regions, see *The MSTP protocol*, on page 9-2.

Specifying a maximum hop number

Applicable to MSTP only, this global property specifies the maximum number of hops that a spanning tree frame (BPDU) can traverse before it is discarded. The default value is **20**. You can specify a value between 1 and 255. This option only appears on the screen when you set the **Mode** property to **MSTP**.

Managing spanning tree instances

By default, the spanning tree protocol STP is enabled on all of the interfaces of the BIG-IP system. The default spanning tree configuration includes a single spanning tree instance, named **0**. A *spanning tree instance* is a discrete spanning tree for a network. While STP and RSTP allow only one spanning tree instance (instance **0**), MSTP allows you to create multiple spanning tree instances, to manage redundant paths for specific VLANs on the network.

When running MSTP, instances that you create have instance members. An *instance member* is a VLAN that you assign to an instance when you create that instance. You can assign as many or as few members to an instance as you deem necessary. By default, all VLANs on the BIG-IP system are members of instance **0**.

If you create an instance and attempt to add a VLAN that is already a member of another instance, the BIG-IP system deletes the VLAN from the existing instance and adds the VLAN to the new instance.

Each instance name must be a numeric value that you assign when you create the instance.

When you manage a spanning tree instance, you can:

- View a list of instances
- Create an instance (MSTP only)
- Modify instance properties
- Delete an instance or its members (MSTP only)

Viewing a list of spanning tree instances

You can view a list of existing spanning tree instances using the Configuration utility. For STP and RSTP, the only instance in the list is instance **0**. For MSTP, the list shows instance **0**, plus any other instances that you have explicitly created. For information on creating a spanning tree instance, see *Configuring interfaces for spanning tree*, on page 9-16.

When you view a list of instances, you can see the following information for each instance:

- The name of the instance
- The bridge priority number
- The MAC address of the root bridge
- The MAC address of the regional root bridge
- The number of instance members

To view a list of spanning tree instances

On the Main tab of the navigation pane, expand **Network**, and click **Spanning Tree**. This opens the Spanning Tree screen, which lists all existing instances.

Creating a spanning tree instance (MSTP-only)

The STP and RSTP protocols allow only one spanning tree instance, instance **0**, which the BIG-IP system creates automatically when you enable spanning tree. When running STP or RSTP, you can modify the properties of instance **0**, but you cannot create additional instances. For information on modifying the properties of an instance, see *Viewing and modifying a spanning tree instance*, on page 9-14.

When you are running MSTP, however, the MSTP algorithm can explicitly create instances. The reason that you can create instances is that MSTP recognizes VLANs. By creating an instance and assigning one or more VLANs to it, you can control bridge loops and redundant paths within those VLANs.

For example, suppose you have two interfaces. One interface is assigned to VLAN A, while the other interface is assigned to VLANs A and B. If you are using the STP or RSTP protocol, both of which disregard VLANs, the protocol might block traffic for both VLANs, as shown in Figure 9.1.

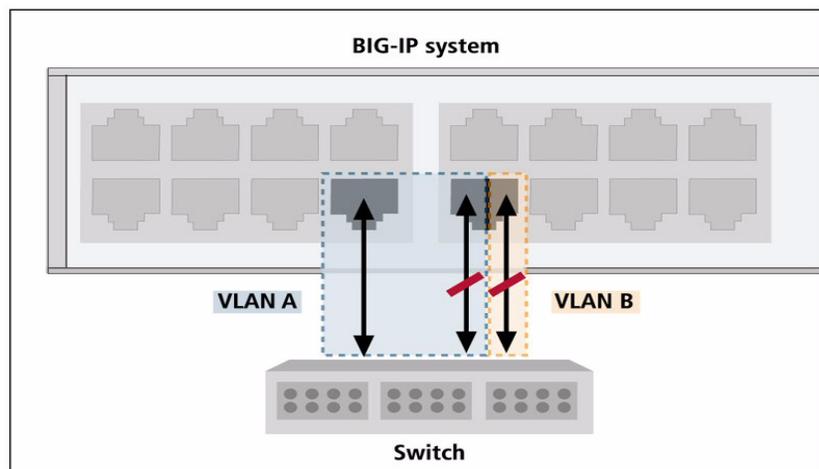


Figure 9.1 Using STP or RSTP to block redundant paths

By contrast, the MSTP protocol can make blocking decisions on a per-VLAN basis. In our example, on the interface that carries traffic for two VLANs, you can block traffic for VLAN A, but leave a path open for VLAN B traffic. This is shown in Figure 9.2, on page 9-12.

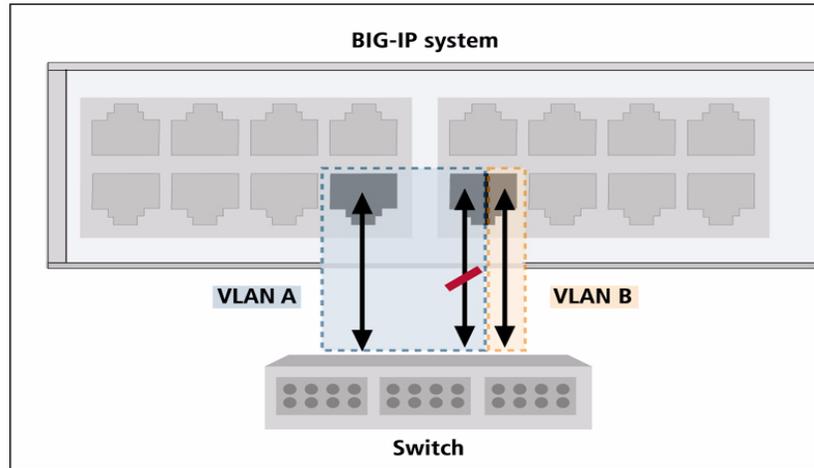


Figure 9.2 Using MSTP to block redundant paths

Table 9.3 shows the properties that you configure when you create or modify a spanning tree instance.

Property	Description	Default Value
Instance ID	Specifies a numeric identification for the instance. This number can be between 1 and 255.	No default value
Bridge Priority	Specifies the spanning tree bridge priority for the instance.	61440
VLANs	For MSTP only, specifies the VLANs that you want to be members of the instance.	external and internal

Table 9.3 Configurable properties of a spanning tree instance

Use the following procedures to create a spanning tree instance. For more information on each property that you configure, see the sections that follow the procedures.

◆ Important

*From the Configuration utility screen that lists existing spanning tree instances, you create instances using a special **Create** button. This **Create** button appears only when you are running MSTP. If you are running MSTP, but no **Create** button appears on the screen, your BIG-IP hardware platform does not support MSTP. For more information on creating a spanning tree instance, see **To create a spanning tree instance (MSTP only)**, following.*

To create a spanning tree instance (MSTP only)

1. On the Main tab of the navigation pane, expand **Network**, and click **Spanning Tree**.
This opens the Spanning Tree screen, which lists all existing spanning tree instances.
2. In the upper right corner of the screen, click **Create**.
This displays the properties for that instance.
*Note: If you do not see a **Create** button, then either the global **Mode** property is not set to **MSTP**, or your **BIG-IP** system hardware platform does not support **MSTP**.*
3. In the **Instance ID** box, type an instance identification number.
4. From the **Bridge Priority** list, select a bridge priority or retain the default value.
For more information, see *Selecting a bridge priority*, on page 9-13.
5. For the **VLANs** property, use the Move button (<<) to add members to the instance, or retain the default members.
*Note: If no VLANs appear in the **Available** box, or you need more information, see **Adding VLANs to an instance**, on page 9-14.*
6. Click **Finished**.

◆ Note

Once you create a spanning tree instance, it is automatically enabled. You do not need to explicitly enable it.

Assigning an instance ID

When you configure the **Instance ID** setting, you specify a numeric value for the instance, in the range of 1 to 255. The reason that instance names must be numeric is to handle the requirement that all cooperating bridges agree on the assignment of VLANs to instance IDs. Using numeric values instead of names makes this requirement easier to manage.

Selecting a bridge priority

The bridge in the spanning tree with the lowest relative priority becomes the root bridge. A **root bridge** represents the root of a spanning tree, and is responsible for managing loop resolution on the network. We recommend that you configure this setting so that the BIG-IP system never becomes the root bridge. For this reason, the default value for the **Bridge Priority** setting is **61440**, the highest value that you can select. Note that a bridge priority must be in increments of 4096.

Adding VLANs to an instance

If you are running MSTP, you can add members to a spanning tree instance. An *instance member* is a VLAN. You add members to an instance by associating one or more VLANs with the instance. The interfaces or trunks associated with each VLAN automatically become part of the spanning tree corresponding to that instance.

For two or more bridges to operate in the same spanning tree, all of those bridges must be in the same region, and therefore must have the same instance numbers, instance members, and VLAN tags.

For example, if a bridge has instance **1**, with two VLAN members whose tags are **1000** and **2000**, then any other bridges that you want to operate in that spanning tree must also have instance **1** with two VLAN members whose tags are **1000** and **2000**. For more information on MSTP regions, see *The MSTP protocol*, on page 9-2.

A particular VLAN cannot be associated with more than one spanning tree instance. For example, if you have two instances named **0** and **1**, you can only associate VLAN **external** with one of those instances, not both. Therefore, before creating an instance, verify that each VLAN you intend to associate with the instance is not a member of another instance.

◆ Tip

*If no VLANs appear in the **Available** box when creating an instance, it is likely that all VLANs on the BIG-IP system are members of other instances. You can verify this by viewing the members of other instances. For more information, see **Viewing and modifying a spanning tree instance**, following.*

Viewing and modifying a spanning tree instance

Using the Configuration utility, you can view and modify properties of any instance, including instance **0**. If you are running MSTP, you can modify the **Bridge Priority** and **VLANs** properties. If you are running RSTP or STP, you can modify only the **Bridge Priority** property. In no case can you modify the instance ID.

The procedure for viewing and modifying the properties of an instance follows. For information on the **Bridge Priority** and **VLANs** properties, see *Selecting a bridge priority*, on page 9-13 and *Adding VLANs to an instance*, on page 9-14.

To view and modify properties of an instance

1. On the Main tab of the navigation pane, expand **Network**, and click **Spanning Tree**.
This displays the list of spanning tree instances.
2. In the Name column, click an instance number.
This displays the properties for that instance.

3. Make any modifications that are available for the particular spanning tree protocol you are using (STP, RSTP, or MSTP).
4. Click **Update**.

Deleting a spanning tree instance or its members (MSTP-only)

If you are running MSTP, you might have explicitly created some spanning tree instances. If so, you can delete any spanning tree instance except instance **0**.

You can also remove VLAN members from an instance. When you remove a VLAN from an instance, the VLAN automatically becomes a member of instance **0**. (By default, instance **0** includes any VLAN that is not a member of another instance.)

If you remove all members from an instance, the BIG-IP system automatically deletes the instance.

◆ **Note**

*If you are running RSTP or STP, you cannot delete instance **0** or remove members from it.*

To delete a spanning tree instance

1. On the Main tab of the navigation pane, expand **Network**, and click **Spanning Tree**.
This displays the list of spanning tree instances.
2. In the Name column, locate the instance number.
3. To the left of the instance number, check the Select box.
4. Click **Delete**.
A confirmation message appears.
5. Click **Delete**.

To delete members from a spanning tree instance

1. On the Main tab of the navigation pane, expand **Network**, and click **Spanning Tree**.
This displays the list of spanning tree instances.
2. In the Name column, click an instance name.
This displays the properties for that instance.
3. For the **VLANs** property, use the Move button (>>) to delete members from the instance.
4. Click **Update**.

Configuring interfaces for spanning tree

Some of the configuration tasks you perform when managing a spanning tree protocol pertain to BIG-IP system interfaces. The interface-related tasks you perform are:

- Configuring settings on each interface that is to be part of the spanning tree
- Managing interfaces per spanning tree instance

Configuring spanning tree settings on an interface

For each interface on the BIG-IP system, there are several STP-related settings that you can configure. Table 9.4 lists these settings.

Setting	Description	Default Value
STP	Specifies whether the interface can participate in the spanning tree. By default, this setting is enabled on all BIG-IP system interfaces.	Enabled
STP Link Type	Specifies the link type so that STP uses the correct optimizations for the interface. Possible values are p2p , Shared , and Auto .	Auto
STP Edge Port	Specifies, when checked, that the interface connects to an end station instead of another spanning tree bridge.	Enabled
STP Edge Port Detection	Specifies, when checked, that the system automatically determines the edge port status of the interface.	Enabled
STP Protocol Detection	Resets the interface back to using the RSTP or MSTP protocol after a legacy bridge has been removed. You must manually reset an interface whenever a legacy bridge is removed from the network.	No default value

Table 9.4 Configuration settings for an interface

Use the following procedure to configure the spanning tree settings of an individual interface. For detailed information on each setting, see the sections following the procedure.

◆ Note

*There are additional interface properties and settings that apply to specific spanning tree instances only. For more information, see **Managing interfaces for a specific instance**, on page 9-19.*

To configure spanning tree settings for an interface

1. On the Main tab of the navigation pane, expand **Network**, and click **Interfaces**.

This displays the list of the interfaces on the BIG-IP system.

2. In the Name column, click an interface name.
This displays the general properties of that interface, as well as some configuration settings.
3. In the STP Configuration area, configure the settings as needed.
For information on each setting, see the following sections.
4. Click **Update**.

Enabling and disabling Spanning Tree

When you check the box for the **STP** setting, you are specifying that the interface can become part of a spanning tree. Once the interface becomes part of the spanning tree, the spanning tree protocol takes control of all learning and frame forwarding on that interface.

If you disable this setting, the spanning tree protocol treats the interface as non-existent, and does not send BPDUs to that interface. Also, the interface, and not the spanning tree protocol, controls all learning and frame forwarding for that interface.

Note that you can also enable or disable spanning tree for a trunk. If spanning tree is enabled on the reference link of a trunk (that is, the lowest-numbered interface of the trunk), then spanning tree is automatically enabled on that trunk. To disable spanning tree for a trunk, simply disable spanning tree on the reference link.

Specifying the STP link type

When you specify an STP link type, you ensure that STP uses the correct optimizations for the interface. Possible values are:

- ◆ **auto**
When you set the STP link type to **auto**, the BIG-IP system determines the spanning tree link type, which is based on the **Active Duplex** interface property.
- ◆ **p2p**
When you set the STP link type to **p2p**, the BIG-IP system uses the optimizations for point-to-point spanning tree links. Point-to-point links connect two spanning tree bridges only. For example, a point-to-point link might connect a 10 Gigabit link to another bridge. For point-to-point links, the **Active Duplex** property interface should be set to **full**. Note that **p2p** is the only valid STP link type for a trunk.
- ◆ **shared**
When you set the STP link type to **shared**, the BIG-IP system uses the optimizations for shared spanning tree links. Shared links connect two or more spanning tree bridges. For example, a shared link might be a 10 Megabit hub. Note that for shared links, the **Active Duplex** interface property should be set to **half**.

Enabling and disabling an STP edge port

When you enable the **STP Edge Port** setting, you are explicitly designating the interface as an edge port. An *edge port* is an interface that connects to an end station rather than to another spanning tree bridge. The default setting is disabled (not checked).

If you would rather have the system automatically designate the interface as an edge port, you can enable the **STP Edge Port Detection** setting instead, described in the following section.

If you enable (check) the **STP Edge Port** setting and the interface subsequently receives STP, RSTP, or MSTP frames (BPDUs), the system disables the setting automatically, because only non-edge interfaces receive BPDUs.

Enabling and disabling detection of an STP edge port

When you enable the **STP Edge Port Detection** setting, the system determines whether the interface is an edge port, and if so, automatically designates the interface as an edge port. The system determines edge port status by monitoring the interface and verifying that it does not receive any incoming STP, RSTP, or MSTP frames (BPDUs).

If the system determines that the interface is not an edge port, but you enabled the **STP Edge Port** setting to explicitly designate the interface as an edge port, the system removes the edge port designation from the interface. No interface that receives BPDUs from a bridge can have edge port status, despite the values of the **STP Edge Port** and **STP Edge Port Detection** settings.

Resetting the spanning tree protocol

As described in *Configuring global spanning tree properties*, on page 9-5, the spanning tree algorithm automatically detects the presence of legacy STP bridges on the network, and falls back to STP mode when communicating with those bridges. Because legacy STP bridges do not send spanning tree BPDUs periodically in all circumstances, the BIG-IP system cannot detect when a legacy STP bridge has been removed from the network. Therefore, it is necessary to manually notify the BIG-IP system that the algorithm can switch to the RSTP or MSTP protocol again, whenever a legacy bridge has been removed.

You reset an interface using the **Reset** button for the **STP Protocol Detection** setting.

Managing interfaces for a specific instance

When you manage an interface for a specific spanning tree instance, you can:

- View a list of interfaces for an instance
- View instance-specific properties of an interface
- Configure instance-specific settings for an interface

Viewing a list of interface IDs for an instance

Using the Configuration utility, you can view a list of the interface IDs associated with a specific spanning tree instance.

If you are using MSTP, the interface IDs that appear in the list are the interfaces assigned to the VLANs that you specified when you created the instance. If you are using STP or RSTP, the interface IDs in the list are those that the BIG-IP system automatically assigned to instance **0**.

The list of interface IDs also displays the following information for each interface:

- The STP instance ID
- The priority
- The external path cost
- The port role

To view a list of interfaces for an instance

1. On the Main tab of the navigation pane, expand **Network**, and click **Spanning Tree**.
This opens the Spanning Tree screen and displays a list of existing instances.
2. In the Name column, click an instance number.
This displays the properties of that instance.
3. On the menu bar, click **Instance Interfaces**.
This displays a list of the interfaces for that instance.

Viewing instance-specific properties of an interface

Once you have used the previous procedure to view the list of interfaces associated with a particular spanning tree instance, you can view the properties associated with that interface. Some of these properties are those that you configured using the Interfaces screen. Table 9.5 shows the per-instance interface properties that you can view.

Setting	Description	Default Value
Port Role	Indicates the spanning tree role of the interface (port) with regard to the spanning tree instance. The system determines the interface role automatically.	No default value
Port State	Indicates the manner in which the interface (port) processes any frames that are not spanning tree frames. The system determines the interface state automatically.	No default value

Table 9.5 *Viewable properties of a per-instance interface*

The following two sections describe the **Port Role** and **Port State** properties. For information on the other properties shown in Table 9.5, see *Configuring spanning tree settings on an interface*, on page 9-16.

Understanding the port roles

The **Port Role** property of a per-instance interface specifies the interface's role in the spanning tree instance. You cannot specify a value for this property; the BIG-IP system automatically assigns a role to the interface.

The BIG-IP system can assign one of the following roles to an instance interface:

- **Disabled**
The interface has no active role in the spanning tree instance.
- **Root**
The interface provides a path to a root bridge.
- **Alternate**
The interface provides an alternate path to a root bridge, if the root interface is unavailable.
- **Designated**
The interface provides a path away from the root bridge.
- **Backup**
The interface provides an alternate path away from the root bridge, if an interface with a port role of **Designated** is unavailable. The **Backup** role assignment is rare.

Understanding port states

The **Port State** property of an interface specifies the way that the interface processes normal data packets. You cannot specify a value for this property; the BIG-IP system automatically assigns a state to the interface.

An interface can be in one of the following states at any given time:

- **Blocking**
The interface disregards any incoming frames, and does not send any outgoing frames.
- **Forwarding**
The interface passes frames as needed.
- **Learning**
The interface is determining information about MAC addresses, and is not yet forwarding frames.

Configuring instance-specific settings for an interface

There are a few settings that you configure for an interface that only pertain to a specific instance. Table 9.6 lists and describes these settings.

Setting	Description	Default Value
Interface Priority	Specifies the interface's priority in relation to the other interfaces that are members of the spanning tree instance.	128
External Path Cost	Specifies the relative cost of sending traffic through the interface.	Depends on interface speed
Internal Path Cost	Specifies the relative cost of sending traffic through the interface to adjacent bridges within a spanning tree region.	Depends on interface speed

Table 9.6 Configurable settings of an interface for a specific instance

Use the following procedure to configure the settings of an interface for a specific instance. For detailed information on each setting, see the sections following the procedure.

To configure interface settings per instance

1. On the Main tab of the navigation pane, expand **Network**, and click **Spanning Tree**.
This opens the Spanning Tree screen and displays a list of existing instances.
2. In the Name column, click an instance number.
This displays the properties of that instance.
3. On the menu bar, click **Instance Interfaces**.
This displays a list of the interfaces for that instance.

4. In the Name column, click an interface number.
This displays properties and settings for that interface, for the relevant instance.
5. In the Configuration area, configure the settings.
For information on these settings, see the following sections.
6. Click **Update**.

Selecting an interface priority

Each interface has an associated priority within a spanning tree instance. The relative values of the interface priorities affect which interfaces the system chooses to carry network traffic. Using the **Interface Priority** setting, you can select the interface's priority in relation to the other interfaces that are members of the spanning tree instance.

Typically, the system is more likely to select interfaces with lower numeric values to carry network traffic. A priority value that you assign to an interface can be in the range of **0** to **240**, in increments of 16. Thus, the value you assign to an interface can be **0, 16, 32, 64**, and so on, up to **240**.

The default priority for an interface is **128**, the middle of the valid range.

Specifying path cost

Each interface has an associated path cost within a spanning tree instance. The **path cost** represents the relative cost of sending network traffic through that interface. When calculating the spanning tree, the spanning tree algorithm attempts to minimize the total path cost between each point of the tree and the root bridge. By manipulating the path costs of different interfaces, you can steer traffic toward paths that are either faster, more reliable, more economical, or have all of these qualities.

The value of a path cost can be in the range of **1** to **200,000,000**, unless you have legacy STP bridges. In that case, because some legacy implementations support a range of only **1** to **65535**, you should use this more restricted range when setting path costs on interfaces.

The default path cost for an interface is based on the maximum speed of the interface rather than the actual speed, as shown in Table 9.7.

Maximum Interface Speed	Default Path Cost
10 Gb/s	2,000
1 Gb/s	20,000
100 Mb/s	200,000
10 Mb/s	2,000,000

Table 9.7 Default path costs based on interface speeds

For example, an interface that has a maximum speed of 1000 Mb/s (1 Gb/s), but is currently running at a speed of 10 Mb/s, has a default path cost of **20,000**.

Link aggregation does not affect the default path cost. For example, if a trunk has four 1 Gb/s interfaces, the default path cost is **20,000**.

For MSTP, you can set two kinds of path costs, external and internal. For STP and RSTP, you can set an external path cost only.

◆ **External path cost**

The **External Path Cost** setting is used to calculate the cost of sending spanning tree traffic through the interface to reach an adjacent spanning tree region. The spanning tree algorithm tries to minimize the total path cost between each point of the tree and the root bridge. The external path cost applies only to those interfaces (and trunks) that are members of instance **0**.

◆ **Internal path cost**

The **Internal Path Cost** setting allows you to specify the relative cost of sending spanning tree traffic through the interface to adjacent bridges within a spanning tree region. Note that the internal path cost applies only to bridges that support the MSTP mode. The internal path cost applies to those interfaces (and trunks) that are members of any instance, including instance **0**.

To summarize, STP and RSTP use external path costs only, and the costs apply to instance **0** interfaces only. MSTP uses both external and internal path costs, and the internal costs apply to interfaces in all spanning tree instances, including instance **0**.



10

Managing User Accounts

- Introducing user account management
- Managing local user accounts
- Managing remote user accounts

Introducing user account management

An important part of managing the BIG-IP system is creating and managing user accounts for BIG-IP system administrators. By creating user accounts for system administrators, you provide additional layers of security. User accounts ensure that the system:

- Verifies the identity of users logging into the system (authentication)
- Controls user access to system resources (authorization)

To enable user authentication and authorization, you assign passwords and user roles to your user accounts. **Passwords** allow you to authenticate your users when they attempt to log in to the BIG-IP system. **User roles** allow you to control user access to BIG-IP system resources.

User accounts for BIG-IP system administrators can reside either locally on the BIG-IP system, or remotely on another type of authentication server. The types of remote servers that you can use to store BIG-IP system user accounts are:

- Active Directory™ servers
- Lightweight Directory Access Protocol (LDAP) servers
- Remote Authentication Dial-in User Service (RADIUS) servers

If you want your user accounts to reside locally on the BIG-IP system, you must create those user accounts on the BIG-IP system and assign user roles to them. If you want your user accounts to reside remotely on an Active Directory, LDAP, or RADIUS server, you do not use the BIG-IP system to create the accounts. Instead, you use the mechanism provided by the server vendor, and you use BIG-IP system strictly to assign user roles to those remote accounts and to maintain those user role assignments over time.

Understanding user account types

There are two types of user accounts on the BIG-IP system: System maintenance accounts and Web UI accounts.

◆ System maintenance accounts

System maintenance accounts are user accounts that you maintain using the Setup utility. There are two types of system maintenance accounts: the **root** account and the **support** account. System maintenance accounts reside locally on the BIG-IP system and grant full access to BIG-IP system resources. You configure and maintain these accounts using the Setup utility.

◆ Web UI accounts

Web UI accounts are user accounts that you create for other BIG-IP system administrators to use. Web UI accounts can reside either locally on the BIG-IP system, or remotely on a remote authentication server. You create and maintain these accounts using the browser-based Configuration utility. Creating Web UI accounts allows you to assign various user roles to those accounts as a way to control system

administrator access to BIG-IP system resources. A special Web UI account is the **admin** account, which automatically exists on any BIG-IP system. For more information on the **admin** account, see *Configuring the admin account*, on page 10-4. For information on user roles, see *Understanding user roles*, following.

You are not required to have any accounts other than the system maintenance accounts (**root** and **support**) and the **admin** Web UI account, but we recommend that you do so, as a way to intelligently control administrator access to system resources.

The tools you use to create and maintain user accounts vary according to the type of account you are managing. Table 10.1 lists the various user accounts for the BIG-IP system and the tools you use to manage them.

Account Name	Creation/Configuration Tool	Maintenance Tool
The root account	Setup utility	Setup utility
The support account	Setup utility	Setup utility
The admin account	Setup utility	Configuration utility
Other Web UI accounts	Configuration utility	Configuration utility

Table 10.1 Tools for managing user accounts

Understanding user roles

User roles are a means of authorization that allows you to control a user's access to BIG-IP system resources. More specifically, a **user role** defines the types of tasks that a user can perform on the BIG-IP system and the tools that the user can use to perform those tasks. When you create a local or remote user account, you can assign one of four user roles to that account.

Table 10.2 lists and describes the various user roles that you can assign to a user account.

User Role	Description
No Access	When a user role is set to No Access , the user cannot view, modify, or create any configuration information for the BIG-IP system.
Guest	The Guest user role grants read-only access to the user, through the Configuration utility only. A user with this user role has no access to the command-line interface. A user with the Guest role can view configuration information, but cannot create new objects or modify existing ones. Users with this access level do not have access to various Configuration utility elements such as Create buttons, Update buttons, and Delete buttons.
Operator	The Operator user role allows the user to view information and to enable or disable nodes. Users with this user role can access the BIG-IP system through the Configuration utility only.
Administrator	This user role provides the user with full access to all administrative tasks. By default, users with this user role can access the BIG-IP system through the Configuration utility and iControl, but not through the command line interface. However, as an option, you can assign users the ability to also access the BIG-IP system through the command-line interface.

Table 10.2 User roles for user accounts

The BIG-IP system automatically assigns a user role to an account when you create that account. The user role that the system assigns to a user account depends on the type of account:

◆ **root and admin accounts**

The BIG-IP system automatically assigns the **Administrator** user role to the system maintenance **root** account and the Web UI **admin** account. You cannot change this user-role assignment. Thus, any user who successfully logs into the BIG-IP system using the **root** or **admin** account has full access to system resources and can perform all administrative tasks.

◆ **Other Web UI accounts**

The BIG-IP system automatically assigns the **No Access** user role to all Web UI accounts other than the **admin** account. If the user account you are using has the **Administrator** role assigned to it, you can change another account's user role from **No Access** to **Guest**, **Operator**, or **Administrator**. For remote user accounts, if you know that most of your users will need some amount of access to system resources, you can configure the BIG-IP system to use a role other than **No Access** as the default user role.

Managing local user accounts

Managing local user accounts refers to the tasks of creating, viewing, modifying, and deleting local Web UI user accounts on the BIG-IP system, using the browser-based Configuration utility.

The Configuration utility stores local user accounts (including user names, passwords, and user roles) in a local user-account database. When a user logs into the BIG-IP system using one of these locally-stored accounts, the BIG-IP system checks the account to determine the access level assigned to that user account.

You assign a user role to an account at the time that you create the account, or by changing the properties of an existing account.

◆ Important

*Except for users who want to change their own passwords, only users with the role of **Administrator** can manage local user accounts.*

Configuring the admin account

A user account called **admin** resides on every BIG-IP system. Although the BIG-IP system creates this account automatically, you must still assign a password to the account before you can use it. To initially set the password for the **admin** account, you must run the Setup utility. To change its password later, you use the Configuration utility's Users screens.

The **admin** account resides in the local user account database on the BIG-IP system. By default, the BIG-IP system assigns the **Administrator** user role, which gives the user of this account full access to all BIG-IP system resources. You cannot change the user role on this account. For information on user roles, see *Understanding user roles*, on page 10-2.

Creating user accounts

When you create a local user account, you must give the account a name and a password. You must also set the user role, either by retaining the default user role or by assigning a new one. The default user role for local, non-system accounts is **No Access**.

Only users who have been granted the **Administrator** role can create user accounts.

To create a user account

1. On the Main tab of the navigation pane, expand **System**, and click **Users**.
The User List screen opens, displaying a list of all Web UI accounts.
2. In the upper right corner of the screen, click **Create**.
The New User screen opens.
3. In the **User Name** box, type a name for the user account.
4. For the **Authentication** setting, type and confirm a password for the account.
For more information on user account passwords, see *Managing remote user accounts*, on page 10-8.
5. To grant an access level other than **No Access**, use the **Web User Role** setting and select one of these options:
 - **Administrator**
You can optionally check the **Allow Console Access** box, which grants access to the BIG-IP system through the command line interface.
 - **Operator**
 - **Guest**
6. Click **Finished**.

Viewing user accounts

Using the Configuration utility, you can easily display a list of existing local user accounts and view the properties of an individual account. Only users who have been granted the **Administrator** role can view the settings of other user accounts.

To display a list of existing user accounts

1. On the Main tab of the navigation pane, expand **System**, and click **Users**.
The User List screen opens, displaying a list of all Web UI accounts.
2. View the list of user accounts.

To view the properties of a user account

1. On the Main tab of the navigation pane, expand **System**, and click **Users**.
The User List screen opens, displaying a list of all Web UI accounts.
2. In the user-account list, find the user account you want to view and click the account name.
This displays the properties of that user account.

Modifying user accounts

You use the Configuration utility to modify the properties of any existing local user account, other than the **root** account. Only users who have been granted the **Administrator** role can modify user accounts other than their own.

When you modify account properties, you can:

- Change the password
- Change the user role
- Allow console access (that is, using SSH) if the account has a user role of **Administrator**

Users with user roles of **Operator** or **Guest** can change their own passwords only. Users with an **Administrator** role can change their own passwords as well as other users' passwords.

To change properties of a user account other than root

1. On the Main tab of the navigation pane, expand **System**, and click **Users**.
The User List screen opens, displaying a list of all Web UI accounts.
2. In the user-account list, click a user account name.
This displays the properties of that account.
3. Change the password, or choose a new user role for the account, or both. If the user account has the **Administrator** role assigned to it, or you are changing the user role to **Administrator**, you can optionally check the **Allow Console Access** box.
4. Click **Update**.

You can also change some properties of the **root** account. Specifically, you can change the password of the **root** account, and you can enable or disable access to the BIG-IP system through SSH.

To change properties of the root account

1. On the Main tab of the navigation pane, expand **System**, and click **Platform**.
The General screen opens.
2. For the **Root Account** setting, type a new password in the **Password** box, and re-type the new password in the **Confirm** box.
3. If you want to grant SSH access, then for the **SSH Access** setting, check the **Enabled** box, and for the **SSH IP Allow** setting, either:
 - Select *** All Addresses**.
 - Select **Specify Range** and type a range of IP addresses.
4. Click **Update**.

◆ Important

*If you have a redundant system configuration and you change the password on the **admin** account, you must also change the password on the peer unit, to ensure that synchronization of configuration data operates correctly.*

Deleting user accounts

If the account you are using has an **Administrator** user role, you can delete other local user accounts. When you delete a local user account, you remove it permanently from the local user-account database on the BIG-IP system.

◆ Note

*You cannot delete the **admin** user account, nor can you delete the user account with which you are logged in.*

To delete a user account

1. On the Main tab of the navigation pane, expand **System**, and click **Users**.
This opens the User List screen, displaying a list of all Web UI accounts.
2. In the user-account list, locate the name of the account you want to delete and click the Select box to the left of the account name.
3. Click the **Delete** button.
A confirmation box appears.
4. Click **Delete** again.

Managing remote user accounts

When you are using a remote authentication server, you create and store Web UI accounts (including user names and passwords) on that remote server, using the mechanism supplied by that server's vendor. In addition to creating and storing user accounts on the remote server, however, you must also use the BIG-IP system to assign user roles to those accounts, for the purpose of controlling user access to BIG-IP system resources.

You assign user roles to remote accounts using the Configuration utility. The Configuration utility stores user-role information in the BIG-IP system's local user-account database. When a user whose account information is stored remotely logs into the BIG-IP system and is granted authentication, the BIG-IP system then checks its local database to determine the user role that you assigned to that user.

If you do not assign a user role to a remote user account, then the BIG-IP system assigns a default user role. You can specify which user role you want the BIG-IP system to assign as a default user role. For more information, see *To configure the default user role*, on page 10-11.

Important

*Only users with the role of **Administrator** can manage user roles for remote user accounts.*

Specifying a remote user-account database

One of the tasks you can perform with the Configuration utility is to specify the type of remote user-account database that currently stores your remote user accounts. The available database types that you can specify are:

- Active Directory™
- Lightweight Directory Access Protocol (LDAP)
- Remote Authentication Dial-In User Service (RADIUS)

When you specify the type of remote database, you can also configure some database settings. Then, once you have specified and configured the remote user account database, you can assign user roles to your remote user accounts. For more information on user roles, see *Assigning user roles*, on page 10-9.

To specify and configure a remote account database

1. On the Main tab of the navigation pane, expand **System**, and click **Users**.
This opens the User List screen, displaying a list of all Web UI accounts.
2. On the menu bar, click Authentication Source.
This opens the Authentication Source screen.
3. Click **Change**.

4. In the **User Directory** setting, select a remote server type.
The screen expands to show additional configuration settings.
5. Configure all settings as needed.
For more information, see the online help.
6. Click **Finished**.

Assigning user roles

You create remote user accounts using the mechanism provided by the vendor of your remote server. Once you have created remote accounts, you then use the Configuration utility to assign user roles to those accounts. More specifically, you can use the Configuration utility to:

- Explicitly assign a user role to an individual remote account
- Change the user role of an account
- Specify a default user role for accounts that do not have explicit user-role designations

Explicitly assigning a user role

As stated in the previous section, you do not use the Configuration utility to create remote user accounts for the BIG-IP system. However, you can use the Configuration utility to explicitly assign user roles to them.

This task of assigning a user role to a remote account is not required, because the BIG-IP system automatically assigns a default user role to a remote account if you do not explicitly do so. For information on configuring the default user role, see *Configuring the default user role*, on page 10-10.

To explicitly assign a user role

1. On the Main tab of the navigation pane, expand **System**, and click **Users**.
This opens the User List screen, displaying a list of all user accounts.
*Note: This list shows only the remote accounts to which you have explicitly assigned a user role. For more information, see *Changing a user role*, on page 10-10.*
2. In the upper-right corner of the screen, click **Create**.
This displays the New User screen.
3. In the **User Name** box, type the name of the remote user to which you want to assign a user role.
4. For the **Web User Role** setting, select a user role.
If you select **Administrator**, the **Allow Console Access** setting appears.

5. If you selected **Administrator** and want to allow access to the BIG-IP system through the command line interface, click the **Allow Console Access** box.
6. Click **Finished**.

Changing a user role

Sometimes you might want to change the user role that you previously assigned to a remote account. To do so, you must change the properties of that account by clicking the account name on the User List screen. Only those remote user accounts to which you have explicitly assigned a user role (using the Configuration utility) appear in the list of user accounts.

If you did not explicitly assign a user role to an account, the account does not appear in the list of user accounts. In this case, you cannot change the authorization properties of that individual account. For more information, see *To explicitly assign a user role*, on page 10-9.

To change the properties of a remote user account

1. On the Main tab of the navigation pane, expand **System**, and click **Users**.
This opens the User List screen, displaying a list of user accounts to which you explicitly assigned user roles.
2. In the User Name column, click a user name.
This displays the properties for that user account.
3. In the **Web User Role** box, select a user role.
4. Click **Update**.

Configuring the default user role

Sometimes, you might have remote user accounts to which you have not explicitly assigned user roles. (For more information, see *Assigning user roles*, on page 10-9.) Such accounts do not appear in the list of user accounts on the User List screen.

To ensure that these accounts have a user role assigned to them, the BIG-IP system automatically assigns a default user role, to ensure valid user authorization. By default, the user role that the BIG-IP system assigns to these remote accounts is **No Access**. However, you can change the user role that the BIG-IP system uses as the default user role, to **Administrator**, **Operator**, or **Guest**. Then, whenever you create a user account on that remote server and you do not explicitly assign a user role to that account, the BIG-IP system automatically assigns that user role to the account.

To configure the default user role

1. On the Main tab of the navigation pane, expand **System**, and click **Users**.
This opens the User List screen, displaying a list of all Web UI accounts.
2. From the Users menu, choose Remote Access.
This displays the Remote Access screen.
3. In the **Web User Role** setting, select a default user role from the list.
If you select the **Administrator** user role, an optional setting appears for granting console access to the user.
4. If you want to grant the user access to the BIG-IP system through the command-line interface, check the **Allow Console Access** box.
5. Click **Update**.

At any time, you can view the default user role that the BIG-IP system assigns to any remote accounts for which you have not explicitly assigned a user role.

To view the current default user role

1. On the Main tab of the navigation pane, expand **System**, and click **Users**.
This opens the User List screen, displaying a list of all Web UI accounts.
2. From the Users menu, choose Remote Accounts.
The **Web User Role** setting displays the user role that the BIG-IP system currently uses as the default user role for remote accounts.

Deleting an explicit user-role designation

When you use the Configuration utility to delete a remote user account, you are not actually deleting the account from the remote server. Instead, you are removing the explicit user-role designation that you previously assigned the account.

Removing an explicit user-role designation from a remote user account causes the BIG-IP system to assign the default user role to the account.

To delete an explicit user role designation

1. On the Main tab of the navigation pane, expand **System**, and click **Users**.
This opens the User List screen, displaying a list of all Web UI accounts.

2. Locate an account name in the list and click the corresponding Select box.
3. Click **Delete**.
A confirmation page appears.
4. Click **Delete**.



II

Configuring SNMP

- Introducing SNMP administration
- Configuring the SNMP agent
- Working with SNMP MIB files
- Collecting performance data

Introducing SNMP administration

Simple Network Management Protocol (SNMP) is an industry-standard protocol that gives a standard SNMP management system the ability to remotely manage a device on the network. One of the devices that an SNMP management system can manage is a BIG-IP system. The SNMP versions that the BIG-IP system supports are: SNMP v1, SNMP v2c, and SNMP v3. The BIG-IP system implementation of SNMP is based on a well-known SNMP package, Net-SNMP, which was formerly known as UCD-SNMP.

Reviewing an industry-standard SNMP implementation

A standard SNMP implementation consists of an *SNMP manager*, which runs on a management system and makes requests to a device, and an *SNMP agent*, which runs on the managed device and fulfills those requests. SNMP device management is based on the standard management information base (MIB) known as MIB-II, as well as object IDs and MIB files.

- The *MIB* defines the standard objects that you can manage for a device, presenting those objects in a hierarchical, tree structure.
- Each object defined in the MIB has a unique object ID (OID), written as a series of integers. An *OID* indicates the location of the object within the MIB tree.
- A set of MIB files resides on both the SNMP manager system and the managed device. *MIB files* specify values for the data objects defined in the MIB. This set of MIB files consists of standard SNMP MIB files and enterprise MIB files. *Enterprise* MIB files are those MIB files that pertain to a particular company, such as F5 Networks, Inc.

Typical SNMP tasks that an SNMP manager performs include polling for data about a device, receiving notifications from a device about specific events, and modifying writable object data.

Reviewing the BIG-IP system SNMP implementation

To comply with the standard SNMP implementation, the BIG-IP system includes both an SNMP agent, a set of standard SNMP MIB files, and a set of enterprise MIB files (those that are specific to the BIG-IP system). The enterprise MIB files typically reside on both the BIG-IP system and system running the SNMP manager. Fortunately, you can use the browser-based Configuration utility to download the enterprise MIB files to your SNMP manager.

Using the BIG-IP system implementation of SNMP, the SNMP manager can perform these distinct functions:

- Poll for information (such as performance metrics)
- Receive notification of specific events that occur on the BIG-IP system
- Set data for SNMP objects that have a read/write access type

The last item in the list refers to the ability of an SNMP manager system to enable or disable various BIG-IP system objects such as virtual servers and nodes. Specifically, you can use SNMP to:

- Enable or disable a virtual server
- Enable or disable a virtual address
- Enable or disable a node
- Enable or disable a pool member
- Set a node to an **up** or **down** state
- Set a pool member to an **up** or **down** state
- Reset statistical data for all BIG-IP objects

Summarizing SNMP configuration on the BIG-IP system

Before an SNMP manager system can manage a BIG-IP system remotely, you must perform a few configuration tasks on the BIG-IP system, using the BIG-IP system's Configuration utility. After you have performed these configuration tasks, you can use standard SNMP commands on the remote manager system to manage the BIG-IP system.

The configuration tasks you perform are:

◆ **Configuring the SNMP agent**

There are a number of things you can do to configure the SNMP agent on the BIG-IP system. For example, you can allow client access to information that the SNMP agent collects, and you can configure the way that the SNMP agent handles SNMP traps. *Traps* are definitions of unsolicited notification messages that the BIG-IP alert system and the SNMP agent send to the SNMP manager when certain events occur.

◆ **Downloading MIB files**

You can download two sets of MIB files to your remote manager system: the standard SNMP MIB files and the enterprise MIB files.

Configuring the SNMP agent

To configure the SNMP agent on the BIG-IP system, you can use the Configuration utility. Configuring the SNMP agent means performing the following tasks:

- ◆ **Configuring BIG-IP system information**
Specify a system contact name and the location of the BIG-IP system.
- ◆ **Configuring client access to the SNMP agent**
Configure the BIG-IP system to allow access to the SNMP agent from an SNMP manager system.
- ◆ **Controlling access to SNMP data**
Assign access levels to SNMP communities or users, to control access to SNMP data.
- ◆ **Configuring Traps**
Enable or disable traps and specify the destination SNMP manager system for SNMP traps.

An alternative way to configure the SNMP agent is by editing certain BIG-IP system configuration files directly. These files are:

- ◆ **/config/snmp/snmpd.conf**
This file contains most of the configuration information for the SNMP agent, including trap information.
- ◆ **/var/net-snmp/snmpd.conf**
Required for SNMP v3 only, this file contains SNMP user names.
Important: You must stop the **snmpd** service prior to editing this file.
- ◆ **/etc/hosts.allow**
This file contains the IP addresses and netmasks for the manager systems that are allowed access to the BIG-IP system.

◆ **WARNING**

You should attempt to edit these files directly only if you are an advanced BIG-IP system administrator. Also, do not attempt to configure any bigdb™ database keys that correspond to SNMP. Doing so could harm your system.

Configuring BIG-IP system information

You can use the Configuration utility to configure the following information:

- ◆ **Contact Information**
The contact information is a MIB-II simple string variable defined by almost all SNMP boxes. The contact name usually contains a user name, as well as an email address.

◆ **Machine Location**

The machine location is a MIB-II variable that almost all machines support. It is a simple string that defines the location of the machine.

To configure system information

1. On the Main tab of the navigation pane, expand **System**, and click **SNMP**.
This opens the SNMP Agent Configuration screen.
2. In the Global Setup area, fill in the boxes.
For more information, see the online help.
3. Click **Update**.

Configuring client access

An SNMP *client* refers to any system running the SNMP manager software for the purpose of remotely managing the BIG-IP system. To set up client access to the BIG-IP system, you specify the IP or network addresses (with netmask as required) from which the SNMP agent can accept requests. (By default, SNMP is enabled only for the BIG-IP system loopback interface, **127.0.0.1**.)

To allow client access to the SNMP agent

1. On the Main tab of the navigation pane, expand **System**, and click **SNMP**.
This opens the SNMP Agent Configuration screen.
2. For the **Client Allow List Type** setting, select **Host** or **Network**, depending on whether the IP address you specify is a host system or a subnet.
3. Type the following information:
 - In the **Address** box, type an IP address or network address from which the SNMP agent can accept requests.
 - If you selected **Network** in step 3, type the netmask in the **Mask** box.
4. Click the **Add** button to add the host or network address to the list of allowed clients.
5. Click **Update**.

Controlling access to SNMP data

To better control access to SNMP data, you can assign an access level to an SNMP v1 or v2c community, or to an SNMP v3 user.

There is a default access level for communities, and this access level is read-only. This means that you cannot write to an individual data object that has a read/write access type until you change the default read-only access level of the community or user.

The way to modify this default access level is by using the Configuration utility to grant read/write access to either a community (for SNMP v1 and v2c) or a user (SNMP v3), for a given OID.

When you set the access level of a community or user to read/write, and an individual data object has a read-only access type, access to the object remains read-only. In short, the access level or type that is the most secure takes precedence when there is a conflict. Table 11.1 illustrates this point.

If the access type of an object is...	And you set the access level of a community or user to...	Then access to the object is...
Read-only	Read-only	Read-only
	Read/write	Read-only
Read/write	Read-only	Read-only
	Read/write	Read/write

Table 11.1 Access control for SNMP data

To grant community access to SNMP data (v1 or v2c only)

1. On the Main tab of the navigation pane, expand **System**, and click **SNMP**.
This opens the SNMP Agent Configuration screen.
2. From the Agent menu, choose Access (v1, v2c).
This displays the SNMP Access screen.
3. In the upper-right corner of the screen, click **Create**.
This displays the New Access Record screen.
4. Select the type of address to which the access record applies, either **IPv4** or **IPv6**.
5. In the **Community** box, type the name of the SNMP community for which you are assigning an access level (in step 9).
6. In the **Source** box, type the source IP address.
7. In the **OID** box, type the OID for the top-most node of the SNMP tree to which the access applies.

8. For the **Access** setting, select an access level, either **Read Only** or **Read/Write**. (This access level applies to the community name you specified in step 6.)
9. Click **Finished**.

To grant access to SNMP data (v3 only)

1. On the Main tab of the navigation pane, expand **System**, and click **SNMP**.
This opens the SNMP Agent Configuration screen.
2. From the Agent menu, choose Access (v3).
This displays the SNMP Access screen.
3. In the upper-right corner of the screen, click **Create**.
This displays the New Access Record screen.
4. In the **User Name** box, type a user name for which you are assigning an access level (in step 9).
5. For the **Authentication** setting, select a type of authentication to use, and then type and confirm the user's password.
6. For the **Privacy** setting, select a privacy protocol, and then do *one* of the following:
 - Type and confirm the user's password.
 - Click the **Use Authentication Password** box.
7. In the **OID** box, type the object identifier (OID) for the top-most node of the SNMP tree to which the access applies.
8. For the **Access** setting, select an access level, either **Read Only** or **Read/Write**. (This access level applies to the user name that you specified in step 5.)
9. Click **Finished**.

When you use the Configuration utility to assign an access level to a community or user, the utility updates the **snmpd.conf** file, assigning only a single access setting to the community or user. There might be times, however, when you want to configure more sophisticated access control. To do this, you must edit the **/config/snmp/snmpd.conf** file directly, instead of using the Configuration utility.

For example, Figure 11.1 shows a sample **snmpd.conf** file when you use the Configuration utility to grant read/write access to a community.

```
rocommunity public default
rwcommunity public1 127.0.0.1 .1.3.6.1.4.1.3375.2.2.10.1
```

Figure 11.1 Sample access-control assignments in the **snmpd.conf** file

In this example, the string **rocommunity** identifies a community named **public** as having the default read only access level (indicated by the strings **ro** and **default**). This read only access level prevents any allowed SNMP manager in community **public** from modifying a data object, even if the object has an access type of read/write.

The string **rwcommunity** identifies a community named **public1** as having a read/write access level (indicated by the string **rw**). This read/write access level allows any allowed SNMP manager in community **public1** to modify a data object under the tree node **1.2.6.1.4.1.3375.2.2.10.1 (ItnVirtualServ)** on the local host **127.0.0.1**, if that data object has an access type of read/write.

For more information, see the man page for the **snmpd.conf** file.

Configuring traps

On the BIG-IP system, **traps** are definitions of unsolicited notification messages that the BIG-IP alert system and the SNMP agent send to the SNMP manager when certain events occur on the BIG-IP system. Configuring SNMP traps on a BIG-IP system means configuring the way that the BIG-IP system handles traps, as well as setting the destination for notifications that the alert system and the SNMP agent send to an SNMP manager.

Enabling traps for specific events

You can configure the SNMP agent on the BIG-IP system to send, or refrain from sending, notifications when the following events occur:

- The SNMP agent on the BIG-IP system stops or starts. By default, this trap is enabled.
- The BIG-IP system receives an authentication warning, generated when a client system attempts to access the SNMP agent. By default, this trap is disabled.
- The BIG-IP system receives any type of warning. By default, this trap is enabled.

To enable traps for specific events

1. On the Main tab of the navigation pane, expand **System**, and click **SNMP**.
This opens the SNMP Agent Configuration screen.
2. From the Traps menu, choose Configuration.
This displays the SNMP Trap Configuration screen.
3. To send traps when someone starts or stops the SNMP agent, verify that the **Agent Start/Stop** box is checked.
4. To send notifications when authentication warnings occur, click the **Agent Authentication** box.

5. To send notifications when certain warnings occur, verify that the **Device** box is checked.
6. Click **Update**.

Setting the trap destination

In addition to enabling certain traps for certain events, you must specify the destination SNMP manager to which the BIG-IP system should send notifications. For SNMP versions 1 and 2c only, you specify a destination system by providing the community name to which the BIG-IP system belongs, the IP address of the SNMP manager, and the target port number of the SNMP manager.

Important

If you are using SNMP V3 and want to configure a trap destination, you do not use the SNMP screens within the Configuration utility. Instead, you configure the `snmpd.conf` file. For more information, see the man page for the `snmpd.conf` file.

To specify a trap destination

1. On the Main tab of the navigation pane, expand **System**, and click **SNMP**.
This opens the SNMP Agent Configuration screen.
2. From the Traps menu, choose Destination.
This displays the SNMP Destination screen.
3. In the upper-right corner of the screen, click **Create**.
This displays the New Trap Record screen.
4. For the **Version** setting, select an SNMP version number.
5. In the **Community** box, type the community name for the SNMP agent running on the BIG-IP system.
6. In the **Destination** box, type the IP address of the SNMP management system.
7. In the **Port** box, type the SNMP management system port number that is to receive the traps.
8. Click **Finished**.

Working with SNMP MIB files

As described earlier, *MIB files* define the SNMP data objects contained in the SNMP MIB. There are two sets of MIB files that typically reside on the BIG-IP system and the SNMP manager system: enterprise MIB files (that is, F5-specific MIB files) and standard SNMP MIB files.

Both sets of MIB files are already present on the BIG-IP system, in the directory `/usr/share/snmp/mibs`. However, you still need to download them to your SNMP manager system. You can download these MIB files from the Welcome screen of the browser-based Configuration utility.

Downloading SNMP MIB files

The enterprise MIB files that you can download to the SNMP manager system are:

- ◆ **F5-BIGIP-COMMON-MIB.txt**

This MIB file contains common information and all notifications (traps). For more information, see *Using the F5-BIGIP-COMMON-MIB.txt file*, on page 11-10.

- ◆ **F5-BIGIP-LOCAL-MIB.txt**

This is an enterprise MIB file that contains specific information for properties associated with specific BIG-IP system features related to local traffic management (such as virtual servers, pools, and SNATs). For more information, see *Using the F5-BIGIP-LOCAL-MIB.txt file*, on page 11-10.

- ◆ **F5-BIGIP-SYSTEM-MIB.txt**

The **F5-BIGIP-SYSTEM-MIB.txt** MIB file includes global information on system-specific objects. For more information, see *Using the F5-BIGIP-SYSTEM-MIB.txt file*, on page 11-11.

To view the set of standard SNMP MIB files that you can download to the SNMP manager system, list the contents of the BIG-IP system directory `/usr/share/snmp/mibs`.

To download MIB files

1. On the Main tab of the navigation pane, expand **Overview**, and click **Welcome**.
This opens the Welcome screen.
2. Scroll to the Downloads section, and locate the **SNMP MIBs** section.
3. Click the type of MIB files to download.
The two MIB file types are F5 MIB files and Net-SNMP MIB files.
4. Follow the instructions on the screen to complete the download.

Understanding the enterprise MIB files

Once you have downloaded all of the necessary MIB files, you should familiarize yourself with the contents of the enterprise MIBs, for purposes of managing the BIG-IP system and troubleshooting BIG-IP system events.

◆ Note

To manage a BIG-IP system with SNMP, you need to use the standard set of SNMP commands. For information on SNMP commands, consult your favorite third-party SNMP documentation, or visit the web site <http://net-snmp.sourceforge.net>.

As mentioned in *Downloading SNMP MIB files*, on page 11-9, the BIG-IP system includes a set of enterprise MIB files:

- **F5-BIGIP-COMMON-MIB.txt**
- **F5-BIGIP-LOCAL-MIB.txt**
- **F5-BIGIP-SYSTEM-MIB.txt**

These MIB files contain information that you can use for your remote management station to: poll the SNMP agent for BIG-IP system-specific information, receive BIG-IP system-specific notifications, or set BIG-IP system data.

Using the F5-BIGIP-COMMON-MIB.txt file

The **F5-BIGIP-COMMON-MIB.txt** file is an enterprise MIB file that contains objects pertaining to any common information, as well as the F5-specific SNMP traps.

All F5-specific traps are contained within this MIB file. You can identify the traps within this MIB file by viewing the file and finding object names that show the designation **NOTIFICATION-TYPE**.

When an F5-specific trap sends a notification to the SNMP manager system, the SNMP manager system receives a text message describing the event or problem that has occurred. For troubleshooting assistance regarding F5-specific traps, see Appendix A, *Troubleshooting SNMP Traps*.

To see all available MIB objects in this MIB file, you can view the **F5-BIGIP-COMMON-MIB.txt** file in the directory `/usr/share/snmp/mibs` on the BIG-IP system.

Using the F5-BIGIP-LOCAL-MIB.txt file

The **F5-BIGIP-LOCAL-MIB.txt** file is an enterprise MIB file that contains information that an SNMP manager system can access for the purpose of managing local application traffic. For example, you can:

- View the maximum number of entries that a node can have open at any given time.
- Get a pool name.
- View the current active members for a load balancing pool.
- Reset pool statistics
- Get profile information such as the total number of concurrent authentication sessions.

In general, you can use this MIB file to get information on any local traffic management object (virtual servers, pools, nodes, profiles, SNATs, health monitors, and iRules). You can also reset statistics for any of these objects.

To see all available enterprise MIB objects for local traffic management, you can view the **F5-BIGIP-LOCAL-MIB.txt** file in the directory **/usr/share/snmp/mibs** on the BIG-IP system.

Using the F5-BIGIP-SYSTEM-MIB.txt file

The **F5-BIGIP-SYSTEM-MIB.txt** file is an enterprise MIB file that describes objects representing common BIG-IP system information. Examples of information in this MIB file are global statistic data, network information, and platform information. Some of the data in this MIB file is similar to that defined in MIB-II, but is not exactly the same.

Table 11.2 shows standard MIB-II objects and the F5-specific objects that approximately correspond to them.

MIB-II Category or Object	F5-BIGIP-SYSTEM-MIB Object Name
MIB-II	f5.bigipSystem
interfaces	sysNetwork.sysInterfaces.sysInterface sysNetwork.sysInterfaces.sysInterfaceStat sysNetwork.sysInterfaces.sysInterfaceMediaOptions
ip	sysGlobalStats.sysGlobalIpStat
ip.AddrTable	sysNetwork.sysSelfIp
ip.RouteTable	sysNetwork.sysRoute
ip.ipNetToMediaTable	sysNetwork.sysArpNdp
icmp	sysGlobalStats.sysGlobalIcmpStat
tcp	sysGlobalStats.sysGlobalTcpStat
udp	sysGlobalStats.sysGlobalUdpStat

Table 11.2 F5-BIGIP-SYSTEM-MIB objects and their relationship to MIB-II objects

MIB-II Category or Object	F5-BIGIP-SYSTEM-MIB Object Name
transmission/dot3.dot3StatTable transmission/dot3.dot3CollTable	sysNetwork.sysTransmission.sysDot3Stat
dot1dBridge.dot1dBase	sysNetwork.sysDot1dBridge
dot1dBridge.dot1dStp	sysNetwork.sysSpanningTree.sysStpBridgeStat sysNetwork.sysSpanningTree.sysStpBridgeTreeStat sysNetwork.sysSpanningTree.sysInterfaceStat sysNetwork.sysSpanningTree.sysInterfaceTreeStat
dot1dBridge.dot1dTp	sysGlobalAttr.VlanFDBTimeout
dot1dBridge.dot1dTpFdbTable	sysNetwork.sysL2
dot1dTpPortTable	sysNetwork.sysInterfaces.sysInterfaceStat
dot1dStaticTable	Not supported.
ifMIB/ifMIBObjects.ifXTable	sysNetwork.sysInterfaces.sysIfxStat

Table 11.2 F5-BIGIP-SYSTEM-MIB objects and their relationship to MIB-II objects

To see all available enterprise MIB system objects, you can view the **F5-BIGIP-SYSTEM-MIB.txt** file in the directory `/usr/share/snmp/mibs` on the BIG-IP system.

Using the RMON-MIB.txt file

One of the MIB files that the BIG-IP system provides is the Remote network Monitoring (RMON) MIB file, **RMON-MIB.txt**. This file is the standard RMON MIB file. However, the implementation of RMON on the BIG-IP system differs slightly from the standard RMON implementation, in these ways:

- The BIG-IP system implementation of RMON supports four of the nine RMON groups. The four supported RMON groups are: statistics, history, alarms, and events.
- The **RMON-MIB.txt** file monitors the BIG-IP system interfaces (that is, **sysIfIndex**), and not the standard LINUX interfaces.
- For hardware reasons, the packet-length-specific statistics in the RMON statistics group offer combined transmission and receiving statistics only. This behavior differs from the behavior described in the definitions of the corresponding object IDs.

To understand how RMON operates for a BIG-IP system, you can view the **RMON-MIB.txt** file in the directory `/usr/share/snmp/mibs` on the BIG-IP system.

Collecting performance data

The Configuration utility on the BIG-IP system displays graphs showing performance metrics for the system. However, you can also use SNMP to collect the same information.

The types of performance metrics that you can gather using SNMP are:

- Megabytes of memory being used
- Number of active connections
- Number of new connections
- Throughput in bits per second
- Number of HTTP requests
- CPU use

Each type of metric has one or more SNMP object IDs (OIDs) associated with it. To gather performance data, you specify these OIDs with the appropriate SNMP command.

For example, the following SNMP command collects data on current memory use, where **public** is the community name and **bigip** is the host name of the BIG-IP system:

```
snmpget -c public bigip sysGlobalStat.sysStatMemoryUsed.0
```

For some types of metrics, such as memory use, simply issuing an SNMP command with an OID gives you the information you need. For other types of metrics, the data that you collect with SNMP is not useful until you perform a calculation on it.

For example, to determine the throughput rate of client bits coming into the BIG-IP system, you must perform the following calculation on the data that you collect with the OID shown:

```
( sysStatClientBytesIn (.1.3.6.1.4.1.3375.2.1.1.2.1.3)*8 ) / time
```

This calculation takes the data resulting from specifying the OID **sysStatClientBytesIn**, multiplies the value by **8**, and divides it by the elapsed time.

The following sections contain tables that list:

- The performance graphs that the Configuration utility displays
- The OIDs that you can use to collect the performance data shown in the graphs
- The calculations that you must perform to interpret the performance data that you collect

◆ Note

If an OID that is listed in any of the following sections does not show a calculation, then no calculation is required.

Collecting data on memory use

You can use an SNMP command with an OID to gather data on the number of megabytes of memory currently being used on the BIG-IP system. Table 11.3 shows the OID that you need to specify to gather data on the current memory use. In this case, you do not need to perform a calculation on the collected data.

Performance Graph (Configuration utility)	Graph Metric	Required SNMP OID
Memory Used	memory usage	sysStatMemoryUsed (.1.3.6.1.4.1.3375.2.1.1.2.1.45)

Table 11.3 Required OIDs for collecting metrics on memory use

Collecting data on active connections

You can use SNMP commands with various OIDs to gather data on the number of active connections on the BIG-IP system. Table 11.4 shows the OIDs that you need to specify to gather data on active connections. In this case, you do not need to perform any calculations on the collected data.

Performance Graph (Configuration utility)	Graph Metrics	Required SNMP OIDs
Active Connections (summary graph)	Connections	sysStatClientCurConns (.1.3.6.1.4.1.3375.2.1.1.2.1.8)
Active Connections (detailed graph)	client	sysStatClientCurConns (.1.3.6.1.4.1.3375.2.1.1.2.1.8)
	server	sysStatServerCurConns (.1.3.6.1.4.1.3375.2.1.1.2.1.15)
	pva client	sysStatPvaClientCurConns (.1.3.6.1.4.1.3375.2.1.1.2.1.22)
	pva server	sysStatPvaServerCurConns (.1.3.6.1.4.1.3375.2.1.1.2.1.29)
	ssl client	sysClientsslStatCurConns (.1.3.6.1.4.1.3375.2.1.1.2.9.2)
	ssl server	sysServersslStatCurConns (.1.3.6.1.4.1.3375.2.1.1.2.10.2)

Table 11.4 Required OIDs for collecting metrics on active connections

Collecting data on new connections

You can use SNMP commands with various OIDs to gather data on the number of new connections on the BIG-IP system. Table 11.5 shows the OIDs that you need to specify to gather data on new connections, along with the calculations that you must perform on the collected data.

Performance Graph (Configuration utility)	Graph Metrics	Required SNMP OIDs and the required calculations
New Connections (summary graph)	Client Connections	sysStatClientTotConns (.1.3.6.1.4.1.3375.2.1.1.2.1.7)
	Client Accepts	sysTcpStatAccepts (.1.3.6.1.4.1.3375.2.1.1.2.12.6) / time
	Server Connects	sysTcpStatConnects (.1.3.6.1.4.1.3375.2.1.1.2.12.8) / time
Total New Connections (detailed graph)	Client Connections	sysStatClientTotConns (.1.3.6.1.4.1.3375.2.1.1.2.1.7) / time
	Server Connections	sysStatServerTotConns (.1.3.6.1.4.1.3375.2.1.1.2.1.14) / time
New PVA Connections (detailed graph)	pva client	sysStatPvaClientTotConns (.1.3.6.1.4.1.3375.2.1.1.2.1.21) / time
	pva server	sysStatPvaServerTotConns (.1.3.6.1.4.1.3375.2.1.1.2.1.28) / time
New SSL Connections (detailed graph)	SSL Client	(sysClientsslStatTotNativeConns (.1.3.6.1.4.1.3375.2.1.1.2.9.6) + sysClientsslStatTotCompatConns (.1.3.6.1.4.1.3375.2.1.1.2.9.9)) / time
	SSL Server	(sysServersslStatTotNativeConns (.1.3.6.1.4.1.3375.2.1.1.2.10.6) + sysServersslStatTotCompatConns (.1.3.6.1.4.1.3375.2.1.1.2.10.9)) / time
New Accepts/Connects (detailed graph)	Client Accepts	sysTcpStatAccepts (.1.3.6.1.4.1.3375.2.1.1.2.12.6) / time
	Server Connects	sysTcpStatConnects (.1.3.6.1.4.1.3375.2.1.1.2.12.8) / time

Table 11.5 Required OIDs for collecting metrics on new connections

Collecting data on throughput

You can use SNMP commands with various OIDs to gather data on the throughput rate on the BIG-IP system, in terms of bits per second. Table 11.6 shows the OIDs that you need to specify to gather data on throughput rate, along with the calculations that you must perform on the collected data.

Performance Graph (Configuration utility)	Graph Metrics	Required SNMP OIDs and the required calculations
Throughput (summary graph)	Client Bits	$(\text{sysStatClientBytesIn} (.1.3.6.1.4.1.3375.2.1.1.2.1.3) + \text{sysStatClientBytesOut} (.1.3.6.1.4.1.3375.2.1.1.2.1.5)) * 8 / \text{time}$
	Server Bits	$(\text{sysStatServerBytesIn} (.1.3.6.1.4.1.3375.2.1.1.2.1.10) + \text{sysStatServerBytesOut} (.1.3.6.1.4.1.3375.2.1.1.2.1.12)) * 8 / \text{time}$
Throughput (detailed graph)	Client Bits In	$(\text{sysStatClientBytesIn} (.1.3.6.1.4.1.3375.2.1.1.2.1.3)) * 8 / \text{time}$
	Client Bits Out	$(\text{sysStatClientBytesOut} (.1.3.6.1.4.1.3375.2.1.1.2.1.5)) * 8 / \text{time}$
	Server Bits In	$(\text{sysStatServerBytesIn} (.1.3.6.1.4.1.3375.2.1.1.2.1.10)) * 8 / \text{time}$
	Server Bits Out	$(\text{sysStatServerBytesOut} (.1.3.6.1.4.1.3375.2.1.1.2.1.12)) * 8 / \text{time}$

Table 11.6 Required OIDs for collecting metrics on throughput

Collecting data on HTTP requests

You can use SNMP commands with various OIDs to gather data on the number of current HTTP requests on the BIG-IP system, in terms of requests per second. Table 11.7 shows the OID that you need to specify to gather data on HTTP requests, along with the calculations that you must perform on the collected data.

Performance Graph (Configuration utility)	Graph Metric	Required SNMP OID and the required calculation
HTTP Requests	HTTP Requests	$\text{sysStatHttpRequests} (.1.3.6.1.4.1.3375.2.1.1.2.1.56) / \text{time}$

Table 11.7 Required OIDs for collecting metrics on HTTP requests

Collecting data on CPU use

You can use SNMP commands with various OIDs to gather data on CPU use on the BIG-IP system. The procedure for gathering data on CPU use is as follows:

1. Perform two separate polls of total cycles (**sysStatTmTotalCycles - .1.3.6.1.4.1.3375.2.1.1.2.1.41**), idle cycles (**sysStatTmIdleCycles - .1.3.6.1.4.1.3375.2.1.1.2.1.42**), and sleep cycles (**sysStatTmSleepCycles - .1.3.6.1.4.1.3375.2.1.1.2.1.43**). Perform these polls approximately ten seconds apart from one another.

2. For each cycle type (total, idle, and sleep), calculate the delta of the values from the two polls. For example:

```
DeltaTmTotalCycles = sysStatTmTotalCycles2 -
sysStatTmTotalCycles1
```

```
DeltaTmIdleCycles = sysStatTmIdleCycles2 -
sysStatTmIdleCycles1
```

```
DeltaTmSleepCycles = sysStatTmSleepCycles2 -
sysStatTmSleepCycles1
```

3. Using the three resulting delta values, use the formula in table 11.8 to calculate the CPU use for that time period.

Performance Graph (Configuration utility)	Graph Metric	Required SNMP OIDs and the required calculation
CPU Usage	TMM CPU Usage	$(\text{DeltaTmTotalCycles} - (\text{DeltaTmIdleCycles} + \text{DeltaTmSleepCycles}) / \text{DeltaTmTotalCycles}) * 100$

Table 11.8 Required OIDs for collecting metrics on CPU use



12

Setting up a Redundant System

- Introducing redundant systems
- Configuring units of a redundant pair
- Specifying the default route on back-end servers
- Synchronizing configuration data
- Continuing with active-active system configuration
- Configuring fail-safe
- Mirroring connection information
- Setting a shared MAC masquerade address
- Maintaining a redundant system

Introducing redundant systems

A **redundant system** is a type of BIG-IP system configuration that allows traffic processing to continue in the event that a BIG-IP system becomes unavailable. A BIG-IP redundant system consists of two identically-configured BIG-IP units. When an event occurs that prevents one of the BIG-IP units from processing network traffic, the peer unit in the redundant system immediately begins processing that traffic, and users experience no interruption in service.

You can configure the units of a redundant system to run in one of two redundancy modes: active/standby or active-active.

- ◆ **Active/standby mode**

With **active/standby** mode, only one of the two units is in an active state, that is, processing traffic, at any given time. The inactive unit serves strictly as a standby unit, becoming active only if the active unit becomes unavailable. When a standby unit becomes active, it normally remains active until an event occurs that requires the other unit to become active again, or until you specifically force it into a standby state. For more information, see *Understanding failover in active/standby mode*, on page 12-3, and *Configuring the redundancy mode*, on page 12-13.

- ◆ **Active-active mode**

With **active-active** mode, both units are in an active state simultaneously, each processing traffic for different virtual servers or SNATs. If an event prevents one of the units from processing traffic, the other unit begins processing that traffic in addition to its own. When the failed unit becomes active again, it does not resume processing connections until you specifically direct it to do so. For more information, see *Understanding failover in active-active mode*, on page 12-4, and *Configuring the redundancy mode*, on page 12-13.

Summary of redundant system features

When you set up a redundant system, you configure a variety of features. Some of these features are required, while others are strictly optional. Table 12.1 shows the complete set of features that you can configure for a redundant system.

Feature	Description
Primary and secondary failover addresses	Not only can you specify the primary self IP addresses that a BIG-IP unit is to use for failover, but you can also specify a second pair of IP addresses, on a different VLAN. This secondary pair of IP addresses is for use in case the primary addresses are unavailable. Primary failover addresses are required, and secondary ones are optional.
Active/standby or active-active configuration	You can configure one unit of the redundant system to be active and the other to remain idle until failover occurs. Or, if you want to use both units of your redundant system to process connections simultaneously, you can configure an active-active system. This feature is required.
Redundancy state preference	You can set up one unit in a pair to be the dominant active BIG-IP system. The unit you set up as the dominant BIG-IP system always attempts to be active. This feature applies to active/standby systems only, and is optional.
Network-based failover	You can configure the BIG-IP system to use the network instead of a hard-wired connection to determine the status of the active unit. This feature is optional.
Configuration synchronization	With this feature, you can configure one BIG-IP unit and then synchronize the configuration with the other BIG-IP unit. This feature is required.
Global display of synchronization status	With this feature, you can display the current synchronization status of a unit on every screen of the Configuration utility. This feature is optional.
System fail-safe	With this feature, the BIG-IP system can monitor various hardware components and system services to detect failures. This feature is optional.
Gateway fail-safe	With this feature, the BIG-IP system can respond to the status of upstream routers. This feature is optional.
VLAN fail-safe	With this feature, a BIG-IP system can take action if a VLAN is no longer able to send or receive traffic. This feature is optional.
Connection mirroring	You can mirror connection and persistence information between redundant units. This enables you to provide seamless failover of client connections. This feature is optional.
Sharing a MAC masquerade address	You can share the media access control (MAC) masquerade address between BIG-IP units in a redundant system. This feature is useful if you want to use the system in a topology with secure hubs. This feature is optional.

Table 12.1 Configurable features of a redundant system

Understanding failover and failback

Perhaps the most important tasks that a redundant system performs are failover (to maintain availability when a specific BIG-IP system becomes unavailable), and failback (to re-establish normal BIG-IP system processing when a previously-unavailable BIG-IP system becomes available again).

What is failover?

Failover is a process that occurs when one system in a redundant system becomes unavailable, thereby requiring the peer unit to assume the processing of traffic originally targeted for the unavailable unit. To facilitate coordination of the failover process, each unit has a unit ID (1 or 2).

An essential element to making failover successful is a feature called configuration synchronization. Configuration synchronization, or **ConfigSync**, is a process where you replicate one unit's main configuration file on the peer unit. Because data is shared in this way, a unit can process the other unit's traffic when failover occurs.

By default, the way that a BIG-IP unit monitors the status of its peer, in order to detect that failover is required, is through a hard-wired connection between the two BIG-IP units. With proper configuration, however, you can cause each BIG-IP unit to monitor peer status by way of a TCP/IP network connection instead. For more information on types of failover connections, see *Configuring the failover type*, on page 12-15.

Understanding failover in active/standby mode

When a redundant system is in active/standby mode, one unit is active, that is, accepting and processing connections on behalf of the redundant system, while the other unit is idle (that is, in a standby state).

When failover occurs, the standby unit becomes active, and it normally stays active until failover occurs again, or until you force it into a standby state. Forcing the unit into a standby state automatically causes the other system to become active.

For example, you can configure unit 1 to process traffic for virtual servers A and B. The standby unit monitors the active unit, and if communications fail, the standby unit initiates a failover and becomes the active unit. It then begins processing traffic for both virtual servers.

You can see an active/standby configuration, first as it behaves normally, and then after failover has occurred, by viewing Figure 12.1.

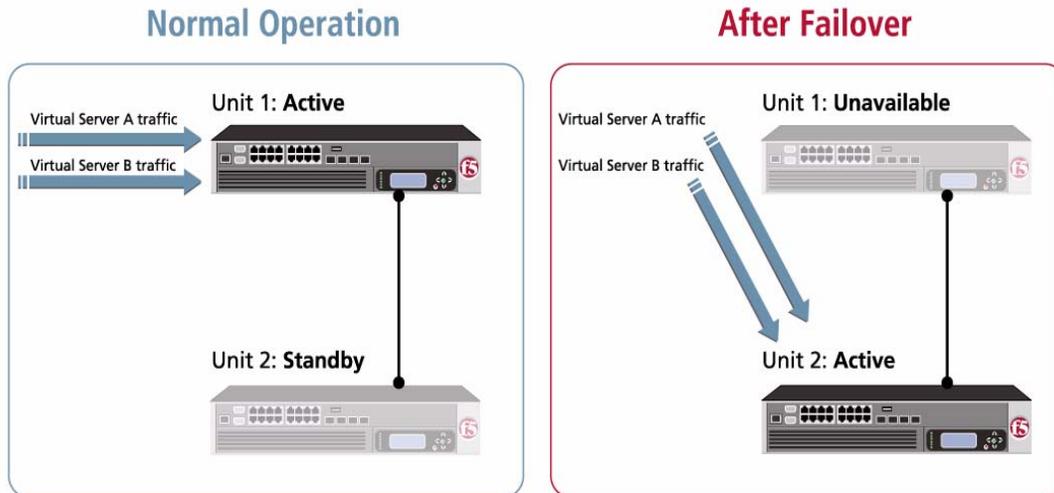


Figure 12.1 Failover on an active/standby system

As you can see in Figure 12.1, unit 1 is in an active state, and unit 2 is in a standby state. With this configuration, failover causes the following to occur:

- Unit 2 switches to an active state.
- Unit 2 begins processing the connections that would normally be processed by its peer.

When the failed unit becomes available again, you can force a unit to change its state from active to standby or from standby to active, thereby initiating failback. **Failback** on an active/standby system causes a unit to relinquish any processing that it is doing on behalf of its peer, and return to a standby state. A redundant system in active/standby mode is the most common type of redundant system. For more information on failback, see *What is failback?*, on page 12-6.

Understanding failover in active-active mode

Unlike an active/standby configuration, which is designed strictly to ensure no interruption of service in the event that a BIG-IP system becomes unavailable, an active-active configuration has an additional benefit. An active-active configuration allows the two units to simultaneously manage traffic, thereby improving overall performance.

A common active-active configuration is one in which each unit processes connections for different virtual servers. For example, you can configure unit 1 to process traffic for virtual servers A and B, and configure unit 2 to process traffic for virtual servers C and D. If unit 1 becomes unavailable, unit 2 begins processing traffic for all four virtual servers.

You can see an active-active configuration, first as it behaves normally, and then after failover has occurred, by viewing Figure 12.2.

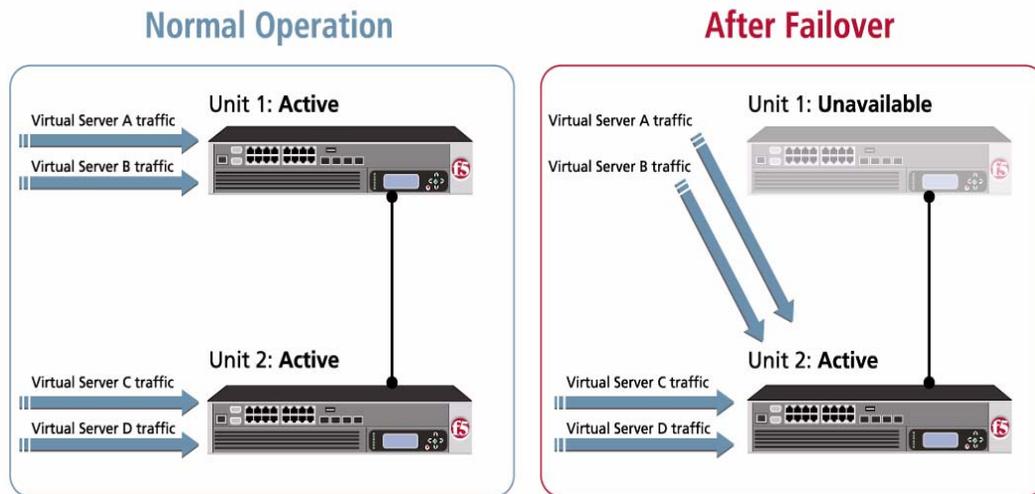


Figure 12.2 Failover on an active-active system

Figure 12.2 shows an active-active configuration in which units 1 and 2 are both in active states. With this configuration, failover causes the following to occur:

- Unit 2 (already in an active state) begins processing the connections that would normally be processed by unit 1.
- Unit 2 continues processing its own connections, in addition to those of unit 1.

When unit 1 becomes available again, you can initiate *failback*, which, in this case, means that the currently-active unit relinquishes any processing that it is doing on behalf of its peer, and continues to operate in an active state, processing its own connections. From this point on, each unit in the active-active system handles its own unit-specific processing. For more information on failback, see *What is failback?*, following.

In addition to associating a unit in an active-active system with a specific virtual server, you can associate a unit with a particular SNAT. Thus, for an active-active configuration to operate successfully, you must associate each virtual server or SNAT with the unit of the redundant system that determines which active unit processes its connections.

What is failback?

Failback is the process of a previously unavailable unit reclaiming its normal traffic as soon as it returns to an active state. Failback behaves differently depending on the mode. For information on managing failback, see *Controlling failback*, on page 12-36.

Understanding self IP addresses for redundant systems

To successfully configure and maintain a redundant system, it is helpful to understand a redundant system's use of static and floating self IP addresses. Configured correctly, static and floating self IP addresses are a key aspect of ensuring that traffic processing continues uninterrupted when failover occurs.

Using static self IP addresses

A **static self IP address** is an IP address that you assign to a BIG-IP system's interface. During normal redundant-system operation prior to failover, the two units use the internal static self IP addresses to continually communicate with one another about system status.

You assign static self IP addresses to each unit's **internal** and **external** VLANs when you run the Setup utility on that unit. Then, when you use the Configuration utility to further configure each unit of your redundant system, you specify a pair of those static self IP addresses, known as a primary failover pair. The **primary** pair specifies the two internal static self IP addresses that the two units use to communicate with one another, before, during, and after failover. (These are the same static self IP addresses that you assigned when you ran the Setup utility on each unit.)

Optionally, you can specify a **secondary** pair, which specifies two alternate static self IP addresses to which failover should occur in the event that the primary addresses are unavailable. Note that the primary addresses typically reside on a different VLAN than the secondary addresses.

For more information on primary and secondary failover addresses, see *Specifying primary and secondary failover addresses*, on page 12-12.

Using floating self IP addresses

A **floating self IP address** is an IP address that is shared between two systems, or between two units of a redundant system. When you run the Setup utility, you normally assign a floating IP address to the internal interface of each unit (in addition to assigning a static self IP address).

Normally, for non-redundant systems (that is, single devices), it is the internal interface's static self IP address that appears as the source IP address in the header of a TCP packet going from a BIG-IP unit to a back-end server. This

causes the back-end server to send its response to that static self IP address. If the BIG-IP system becomes unavailable, the system fails to receive the response.

With a redundant system, however, you can configure the back-end servers to send their responses to a shared floating IP address instead of a static self IP address, and if the target unit is unavailable, the peer unit can receive and process that traffic. Without this shared floating IP address, the delivery of back-end server traffic to the surviving BIG-IP unit would fail.

◆ **Important**

For an active/standby system, when you run Setup on each unit, you must assign the same internal floating IP address to each unit's internal interface. For an active-active system, you must assign a unique floating IP address to each unit's internal interface.

It is sufficient for units of an active/standby configuration to share the same internal floating IP address, but it is not sufficient for an active-active configuration, for these reasons:

- For an active/standby configuration, sharing one internal floating IP address between the two units is sufficient because in the event of failover, the failover unit can always use this address to receive and process all of the unavailable unit's incoming back-end server traffic.
- For an active-active configuration, each unit needs its own internal floating IP address, which the unit shares with its peer (using ConfigSync). Then, based on which of the two floating IP addresses a given back-end server uses to send its responses, the surviving unit can correlate the response to the correct unit's virtual servers and SNATs, and process the response accordingly. For information on configuring back-end servers to use floating IP addresses, see *Specifying the default route on back-end servers*, on page 12-16.

Example of using a floating IP address for an active/standby system

For an active/standby system, suppose that when you initially ran the Setup utility on unit 1, you specified **11.12.11.3** as the internal floating IP address, and when you ran Setup on unit 2, you also specified **11.12.11.3** as its internal floating IP address. When you synchronize the configurations later, **11.12.11.3** should appear on both units as the floating IP address belonging to unit 1.

Then, if unit 1 fails over:

- Unit 2 assumes the unit 1 internal floating IP address (**11.12.11.3**).
- The back-end servers that normally send traffic to the internal address **11.12.11.3** on unit 1 continue to send their traffic to that same address, even though this incoming traffic is now processed by unit 2.
- Unless configured otherwise, unit 2 continues processing traffic until failover occurs again.

Example of using floating IP addresses for an active-active system

For an active-active system, suppose that when you initially ran the Setup utility on unit 1, you specified **11.12.11.3** as the internal floating IP address, and when you ran Setup on unit 2, you specified **11.12.11.4** as its internal floating IP address. When you synchronize the configurations later, **11.12.11.3** should appear on both units as the floating IP address belonging to unit 1, and **11.12.11.4** should appear on both units as the floating IP addresses belonging to unit 2.

Then, if unit 1 fails over:

- Unit 2 assumes the internal floating IP address of unit 1 (**11.12.11.3**).
- The back-end servers that normally send traffic to the internal address **11.12.11.3** on unit 1 continue to send their traffic to that same address, even though this incoming traffic is now processed by unit 2.
- The back-end servers that normally send traffic to the internal address **11.12.11.4** on unit 2 continue to send their traffic to that same address.

Conversely, if unit 2 fails over:

- Unit 1 assumes the internal floating IP address of unit 2 (**11.12.11.4**).
- The back-end servers that normally send traffic to the internal address **11.12.11.4** on unit 2 continue to send their traffic to that same address, even though this incoming traffic is now processed by unit 1.
- The back-end servers that normally send traffic to the internal address **11.12.11.3** on unit 1 continue to send their traffic to that same address.

◆ Tip

*You can configure additional floating IP addresses on the external VLANs of each BIG-IP system as well. This makes it possible for routers to route to a virtual server using virtual **noarp** mode.*

Understanding fail-safe

Fail-safe is the ability of a unit in a redundant system configuration to monitor certain aspects of the system or network, detect interruptions, and consequently take some action, such as initiating failover to the peer unit. Fail-safe can apply to system services, traffic between the BIG-IP system and a gateway router, or VLAN traffic. For more information, see *Configuring fail-safe*, on page 12-23.

Before you begin

Before you begin using this chapter to set up a redundant system, you must run the Setup utility on each unit that is to make up the redundant system if you have not already done so. When you run the Setup utility on a BIG-IP system, you provide some essential information:

- A designation that the system is either a single device or part of a redundant pair
- Static self IP addresses for the VLANs
- Floating self IP addresses for the VLANs

Once you have supplied this information, you can use the remainder of this chapter to complete the configuration of your redundant system. You must first decide, however, whether you want your redundant system to operate in active/standby mode or active-active mode. Then you can use the Configuration utility to configure the two units of your redundant system accordingly. You also need to perform one other configuration task on the back-end servers to which the redundant system sends network traffic.

The remainder of this chapter describes in detail all of the tasks that you must perform to set up a redundant system. In summary, you must:

- ◆ Configure individual settings on each BIG-IP unit, such as redundancy mode and failover IP addresses. For more information, see *Configuring units of a redundant pair*, on page 12-11.
- ◆ Configure the default route on each back-end server. For more information, see *Specifying the default route on back-end servers*, on page 12-16.
- ◆ Synchronize the configuration data from unit 1 to unit 2. For more information, see *Synchronizing configuration data*, on page 12-18.
- ◆ If using active-active mode, associate each virtual server, self IP address, and SNAT with a unit ID, and synchronize the configuration data from unit 2 to unit 1. For more information, see *Continuing with active-active system configuration*, on page 12-21.
- ◆ Configure fail-safe settings for system hardware and services, gateway router traffic, and VLAN traffic. For more information, see *Configuring fail-safe*, on page 12-23.

You can also perform two optional tasks:

- ◆ Mirror connection information. For more information, see *Mirroring connection information*, on page 12-28.
- ◆ Set a shared MAC masquerade address. For more information, see *Setting a shared MAC masquerade address*, on page 12-29.

Once your redundant system is operational, the BIG-IP system displays an indicator in the upper-left corner of all Configuration utility screens, to indicate this information:

- The unit you are currently managing (unit 1 or unit 2)
- The current state of the unit (active or standby)
- Configuration synchronization status (this display is optional)

◆ **Note**

*For information on how to manage a unit of a redundant system on an ongoing basis, see **Maintaining a redundant system**, on page 12-32.*

Configuring units of a redundant pair

To configure a redundant system, you must first configure certain settings on each unit that is to be part of the redundant pair, using the Configuration utility. For some settings, you can simply use a default value.

The settings that you must configure are:

- Primary and secondary failover addresses for a unit and its peer
- The redundancy mode (active/standby or active-active)
- For an active/standby system, the preferred redundancy state for a unit (none, active, or standby)
- The failover method (hard-wired or network)
- The link down time after a failure

Before you configure these settings on a unit, however, you should ensure that the unit is already designated as being part of a redundant pair. (You typically designate a unit as being part of a redundant pair when you initially run the Setup utility on that unit.)

To ensure redundant-system designation for a unit

1. On the Main tab of the navigation pane, expand **System**, and click **Platform**.
The General screen opens.
2. In the General Properties area, verify that the **High Availability** list is set to **Redundant Pair**. If not, select this value.
3. For the **Unit ID** list, retain or change the value, depending on which unit ID you want to assign to this BIG-IP system.
4. Click **Update**.

After verifying the unit's redundant-system designation, use the Configuration utility to perform the following procedure on each unit. When performed on each unit, this procedure creates either an active/standby or an active-active system. Note that the steps in this procedure provide only basic configuration information; the pages following the procedure provide more detailed information for each step, to help you configure the settings in a way that best suits your needs.

Important

*The default type of redundant system that you create using this procedure is an active/standby system. For active-active configurations, you must perform additional configuration tasks after using the Configuration utility. For detailed information, see **Continuing with active-active system configuration**, on page 12-21.*

To configure a redundant system

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. For the **Primary Failover Address** boxes, in the **Self** box type the primary static self IP address for the unit that you are currently configuring, and in the **Peer** box type the primary static self IP address for the peer unit. For more information, see *Specifying primary and secondary failover addresses*, on page 12-12.
Note: Before typing the IP addresses, delete the two colons (::) in each box.
3. If you want to configure the **Secondary Failover Address** settings, in the **Self** box type the secondary static self IP address for the unit you are currently configuring, and in the **Peer** box type the secondary static self IP address for the peer unit. For more information, see *Specifying primary and secondary failover addresses*, on page 12-12.
Note: Before typing the IP addresses, delete the two colons (::) in each box.
4. From the **Redundancy Mode** list, either retain the default value of **Active/Standby**, or select **Active/Active**. For more information, see *Configuring the redundancy mode*, on page 12-13.
5. From the **Redundancy State Preference** list, select a preferred state for the unit you are configuring. For more information, see *Specifying a redundancy state preference*, on page 12-14.
6. If you want the system to use the network type of failover rather than hard-wired failover, check the **Network Failover** box. For more information, see *Configuring the failover type*, on page 12-15.
7. If you want an interface down time other than **0.0**, type a value in the **Link Down Time on Failover** box. For more information, see *Specifying the link down-time on failover*, on page 12-15.
8. Click **Update**.

Specifying primary and secondary failover addresses

When you configure a redundant system with the Configuration utility, you typically specify the internal static self IP addresses of the two peer systems in the redundant configuration. Each IP address that you specify for a unit is the static self IP address that you previously assigned to the unit's internal interface when you ran the Setup utility on that unit.

Known as the *primary failover addresses*, these are the IP addresses that the two BIG-IP units use to communicate with each other before, during, and after failover.

You use the **Self** setting to specify the internal static self IP address of the unit you are currently configuring as part of the redundant system. You use the **Peer** setting to specify the internal static self IP address of the other unit that you want to configure as part of the redundant system.

◆ Tip

*For background information on static self IP addresses, see **Using static self IP addresses**, on page 12-6.*

For example, suppose you want to configure an active/standby system in which the active unit has an internal static self IP address of **10.10.10.1** (unit 1) and the standby unit has an internal static self IP address of **10.10.10.2** (unit 2):

- On unit 1, you set the **Self** and **Peer** addresses in the **Primary Failover Addresses** setting to **10.10.10.1** and **10.10.10.2**, respectively.
- On unit 2, you set the **Self** and **Peer** addresses in the **Primary Failover Addresses** setting to **10.10.10.2** and **10.10.10.1**, respectively.

Optionally, you can specify *secondary failover addresses*, which are static self IP addresses that the BIG-IP system uses if the primary failover addresses are unavailable for some reason.

◆ Note

Secondary failover addresses cannot reside on the same VLAN as primary failover addresses.

◆ Important

*For the **Primary Failover Address** and **Secondary Failover Address** settings, check that you have removed the two colons (:) from these boxes. Failure to do so could cause problems when you synchronize the configuration of the two units.*

Configuring the redundancy mode

Configuring the redundancy mode means specifying whether your redundant system runs in active/standby mode or active-active mode. If you choose active/standby mode, you can specify whether you prefer that unit to be an active or a standby unit when both units are available for processing traffic. For information on setting a preference for a unit's redundancy state, see *Specifying a redundancy state preference*, on page 12-14.

◆ WARNING

MAC masquerading is not supported in active-active mode.

For more information on active/standby and active-active configurations, see *Understanding failover and failback*, on page 12-2. For more information on MAC masquerading, see *Setting a shared MAC masquerade address*, on page 12-29.

Specifying a redundancy state preference

When you configure a system to be part of an active/standby redundant system, you can specify whether you want that system to function primarily as the active system or the standby system in the event that both units can be active at the same time. You can also specify that you have no preference.

The preferences you can set are:

- ◆ **None**
Specifies that this unit does not have a preferred redundancy state. In this case, failback does not normally occur, because a standby unit that becomes active due to failover remains active until failover occurs again. However, you can actively initiate failback when the unavailable unit becomes available again, by forcing the currently-active unit to revert to a standby state.
- ◆ **Active**
Specifies that this unit is the preferred active unit. If you choose this option, the unit can be in a standby state due to a failover event. However, failback to this unit is automatic when this unit becomes available.
- ◆ **Standby**
Specifies that this unit is the preferred standby unit. If you choose this option, then the unit can be in an active state due to a failover condition. However, failback to the peer unit is automatic when the peer unit becomes available.

A redundant system unit that prefers to be active can still serve as the standby unit when the redundant system already has an active unit. For example, if an active unit that prefers to be active fails over and is taken out of service for repair, it can then go back into service as the standby unit until the next time that the redundant system needs an active unit, for example, at reboot.

Configuring the failover type

To enable a unit to fail over to its peer system, you must first specify the type of failover that you want the redundant system to use. The two possible failover types are hard-wired failover and network-based failover.

- **Hard-wired failover**

When you configure *hard-wired* failover, you enable failover by using a failover cable to physically connect the two redundant units. This is the default setting. For the procedure on configuring hard-wired failover, see *Platform Guide: 1500, 3400, 6400, and 6800*.

- **Network failover**

When you configure *network* failover, you enable failover by configuring your redundant system to use the network to determine the status of the active unit. You can use network failover in addition to, or instead of, hard-wired failover.

For background information on failover, see *Understanding failover and failback*, on page 12-2.

Specifying the link down-time on failover

When configuring a unit of a redundant system configuration, you can specify the amount of time that the interfaces are down when the unit fails over to a standby state. Use the **Link Down Time on Failover** setting to prompt peer switches to reset and relearn their ARP tables after a failover.

Specifying the default route on back-end servers

Once you have used the Redundancy Properties screen of the Configuration utility to configure various settings on each unit, you can perform the next task, configuring the default route on each back-end server. (For a high-level summary of tasks, see *Before you begin*, on page 12-8.)

◆ Tip

For any back-end server that receives requests from a unit of a redundant system, you do not need to perform this step if the relevant virtual server is associated with a SNAT.

When failover occurs on a BIG-IP system, the surviving BIG-IP unit begins handling the inbound virtual-server connections (those targeted to back-end servers) that are normally processed by the failed unit, and begins handling the outbound connections (those originating from back-end servers) that are normally destined for the failed BIG-IP unit.

For a redundant system to do this properly, you need to set the default route for each back-end server to the shared, floating IP address assigned to the BIG-IP unit that normally processes the server's responses. This ensures that the back-end server can successfully send a response to the surviving BIG-IP unit.

For example, for an active/standby configuration, if the server **http_server** normally receives connections from unit 1 of your redundant system, and the floating IP address shared by the two units is **11.12.11.3**, you must set the default route for server **http_server** to **11.12.11.3**. Then, if unit 1 goes out of service, the surviving unit (unit 2) can receive the server's response because the IP address **11.12.11.3** is a shared IP address.

◆ Note

Follow your server vendor's instructions for the procedure on setting a default route.

Special considerations for active-active systems

When the surviving unit takes over a virtual server of the failed unit (in addition to having its own virtual servers), the surviving unit must know which virtual server to use to process traffic that it receives from a back-end server. The way it knows this is by the floating IP address to which the response was sent. (As described previously, in an active-active configuration, each unit must have its own unique floating IP address.)

For example, suppose the following:

- The floating IP address of unit 1 is **11.12.11.3**, and the floating IP address of unit 2 is **11.12.11.4**.

- Server **http_server** normally handles traffic for unit 1, which has virtual server **vs_http**. Server **smtp_server** normally handles traffic for unit 2, which has virtual server **vs_smtp**.
- You have configured the default routes of the servers accordingly. That is, you have set the default route of **http_server** to **11.12.11.3**, and the default route of **smtp_server** to **11.12.11.4**.

If unit 1 fails, the surviving unit (unit 2) now has both virtual servers (**vs_http** and **vs_smtp**).

Continuing with our previous example of back-end server **http_server**, if the surviving unit receives a response from **http_server** (sent to **11.12.11.3**, the unit 1 floating IP address), the surviving unit knows, based on that IP address, that the traffic is to be processed by the unit 1 virtual server **vs_http** (now on unit 2).

Similarly, if the surviving unit receives a response from server **smtp_server** (sent to **11.12.11.4**, the unit 2 floating IP address), the surviving unit knows, based on that IP address, that the traffic is to be processed by its own virtual server, **vs_smtp**.

Thus, if you have 20 servers and half of them normally serve the unit 1 virtual servers, you must configure the default route for those 10 servers to be the floating IP address of unit 1. For the remaining 10 servers in your network, which serve the unit 2 virtual servers, you must configure their default route to be the floating IP address of unit 2.

By setting the default routes of your back-end servers to the floating IP address of either unit 1 or unit 2, you ensure that if a unit becomes unavailable, any connections that the servers normally send to that unit are received and processed by the surviving unit, because the surviving unit has assumed the internal floating IP address of the failed machine.

◆ **Tip**

*For information on initially assigning floating IP addresses to redundant-system units, see **Using floating self IP addresses**, on page 12-6.*

The next step

When you have completed the procedure for configuring a redundant system, you need to synchronize the configuration data. To do this, proceed to *Synchronizing configuration data*, on page 12-18.

Note that if you are setting up an active-active system, you must perform additional configuration tasks after performing synchronization. For more information, see *Continuing with active-active system configuration*, on page 12-21.

Synchronizing configuration data

Once you have completed the initial configuration of one of the units in your redundant system, you must synchronize the configuration between the two units. For an active/standby system, you must perform configuration synchronization from the active unit to the standby unit. For an active-active system, you must perform synchronization from each unit to the other.

When you synchronize data from one unit to another, you are giving the target unit the data that it needs to assume traffic processing for its peer when failover occurs. Examples of configuration data that a target unit receives during configuration synchronization are virtual servers, SNATs, and floating IP addresses.

◆ Important

*If you plan to run your redundant system in active-active mode, you must perform other configuration tasks after you have initially synchronized the configuration data from one unit to another. After performing those additional tasks, you must then synchronize the configuration data again. For more information, see **Continuing with active-active system configuration**, on page 12-21.*

You can synchronize configuration data from the current system to the peer system, or you can synchronize the data from the peer system to the current system. The method you choose depends on which unit's data has most recently changed and which unit you are currently configuring.

When you perform configuration synchronization, the BIG-IP system copies a **.ucs** file containing configuration data from one unit to the other. You must synchronize the configuration data before you can put any redundant system into operation.

◆ Note

Prior to synchronizing configuration data, the BIG-IP system creates a backup configuration file on the target system, as a preventative measure.

With respect to configuration synchronization, you can use the Configuration utility to:

- Perform configuration synchronization
- Enable the global display of synchronization status
- View peer synchronization status (for more information, see *Viewing synchronization status*, on page 12-33)

Performing configuration synchronization

When you have a redundant system configuration, it is essential that each unit shares, or *synchronizes*, its current configuration data with its peer unit. If a unit does not share its configuration data with its peer, the surviving unit cannot process traffic for that peer unit. For this reason, you must synchronize configuration data when you initially configure the redundant system, and then repeatedly, on an ongoing basis. The need to repeatedly synchronize data is because a unit's configuration data typically changes over time during normal system maintenance, such as when you add a virtual server or create a new profile, and the unit must share those changes with its peer.

◆ Tip

*You can easily determine when to synchronize configuration data by using the Configuration utility to view synchronization status. Synchronization status indicates whether the configuration data of the units is synchronized. For more information, see **Viewing synchronization status**, on page 12-33.*

You synchronize configuration data using the ConfigSync screen that is available from the Redundant Properties screen of the Configuration utility. The ConfigSync screen contains two settings that are directly related to performing configuration synchronization between redundant units: **ConfigSync User** and **Synchronize**.

The **ConfigSync User** setting allows you to select a user name from a list and type a password. Only a user with appropriate permissions can perform a synchronization. To display this setting, you must select **Advanced** next to the Configuration heading.

◆ Note

*The default value for the **ConfigSync User** setting is **admin**. You do not need to type the password for the **admin** account.*

◆ Important

*The passwords for the **admin** account on each of the BIG-IP units must be the same in order for configuration synchronization to function correctly.*

The **Synchronize** setting allows you to perform a configuration synchronization between two units. To synchronize data, you can click either of the following buttons:

- ◆ **Synchronize TO Peer**

Use this button when the unit you are currently configuring contains updated configuration data that you want to share with the peer unit.

- ◆ **Synchronize FROM Peer**

Use this button when the peer unit contains updated configuration data that you want to share with the unit you are currently configuring.

In either case, the term *peer* refers to the unit with the self IP address that appears in the **ConfigSync Peer** box.

To synchronize configuration data

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. On the menu bar, click **ConfigSync**.
The ConfigSync screen opens.
3. If the user name you are using is other than **admin**:
 - a) Find the Configuration heading and choose **Advanced**.
 - b) For the **ConfigSync User** setting, select a user name.
 - c) In the password box, type a password.
4. For the **Synchronize** setting, click **Synchronize TO Peer** or **Synchronize FROM Peer**.
5. Click **Update**.

Enabling the global display of synchronization status

To ensure that you are aware of any need to synchronize configuration data, you can display synchronization status on each screen of the Configuration utility. You enable this display of synchronization status on the ConfigSync screen, using the **Detect ConfigSync Status** setting. If you check this box, the BIG-IP system displays, on every screen of the Configuration utility, the current synchronization status of the unit you are configuring.

To enable the global display of synchronization status

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. On the menu bar, click ConfigSync.
The ConfigSync screen opens.
3. For the **Detect ConfigSync Status** setting, check the **Enabled** box.
4. Click **Update**.

◆ Tip

*The Configuration utility also displays more detailed synchronization status information. For complete information on viewing configuration synchronization status, see **Viewing synchronization status**, on page 12-33.*

Continuing with active-active system configuration

To fully configure an active-active redundant system, you must perform other tasks in addition to those described in *Before you begin*, on page 12-8. These tasks are:

- Associating the virtual servers and SNATs with the relevant unit number. You perform this task on both units.
- Synchronizing the configuration to the peer unit. You perform this task on both units.

You use the **bigpipe** command-line utility to perform the first task. You use the Configuration utility to perform the second task, synchronizing the configuration.

Associating BIG-IP system objects with unit IDs

Each BIG-IP system in an active-active configuration has a unit ID, either **1** or **2**. When you define a local traffic management object, such as a virtual server, you must associate that object with a specific unit of the active-active redundant pair. When failover occurs, these associations of objects to unit IDs allow the surviving unit to process connections correctly for itself and the failed unit.

You must associate these local traffic management objects with a unit ID:

- Virtual servers
- Self IP addresses
- SNATs

For example, associating virtual server A with unit 1 causes unit 1 to process connections for virtual server A. Associating virtual server B with unit 2 causes unit 2 to process connections for virtual server B. This allows the two units to process traffic for different virtual servers simultaneously, and results in an increase in overall performance. If one of the units fails over, the remaining unit begins processing the connections for all virtual servers of the redundant pair, until failback occurs.

This scenario of using the two units to process different connections simultaneously is one reason for the requirement that both units store identical configuration files (**/config/bigip.conf**).

If you do not associate an object with a specific unit ID in an active-active redundant pair, the redundant system uses **1** as the default unit ID.

Associating a virtual server with a unit ID

You can view a list of virtual servers and their associated unit IDs, and you can change the unit ID associated with a specific virtual server. You can perform these tasks using the **bigpipe** command-line utility.

To view an existing virtual server-unit ID association

To view the unit ID associated with your existing virtual servers, use this **bigpipe** command syntax:

```
b virtual address [<ip addr list> | all] unit [show]
```

To change the unit ID associated with a virtual server

To change the unit ID associated with an existing virtual server, type this command sequence:

```
b virtual address <ip addr> unit <id>
```

Associating a self IP address with a unit ID

You can view a list of self IP addresses and their associated unit IDs, and you can change the unit ID associated with a specific self IP address.

To view an existing self IP address-unit ID association

To view the unit ID associated with your existing self IP addresses, use this **bigpipe** command syntax:

```
b self [<ip addr list> | all] unit [show]
```

To change the unit ID associated with a self IP address

To change the unit ID associated with an existing self IP address, type this command sequence:

```
b self <ip addr> unit <id>
```

Associating a SNAT with a unit ID

You can view a list of SNATS and their associated unit IDs, and you can change the unit ID associated with a specific SNAT. You can perform these tasks using the **bigpipe** command-line utility.

◆ Note

You cannot associate a default SNAT with a unit ID. The default SNAT is not compatible with an active-active system.

To view an existing SNAT-unit ID association

To view the unit ID associated with your existing SNAT translation addresses, use this **bigpipe** command syntax:

```
b snat translation [<ip addr list> | all] unit [show]
```

To change the unit ID associated with a SNAT address

To change the unit ID associated with an existing SNAT translation address, type this command sequence:

```
b snat translation <ip addr> unit <id>
```

Synchronizing the configuration

When you used the Redundancy Properties screen to initially configure each unit of the active-active system, you also synchronized the configurations between the two units. Now that you have associated unit IDs with each virtual server, self IP address, and SNAT, you must synchronize the configurations again. For more information, see *Synchronizing configuration data*, on page 12-18.

Configuring fail-safe

Fail-safe is the ability of a BIG-IP system to monitor certain aspects of the system or network, detect interruptions, and consequently take some action. In the case of a redundant system, a unit can detect a problem and initiate failover to the peer unit. When you configure the fail-safe feature on a BIG-IP system, you are specifying the particular events that cause failover to occur in a redundant system. The fail-safe feature applies to:

- System services
- Traffic between the BIG-IP system and a gateway router
- Traffic on a VLAN

Configuring system fail-safe

When you configure system fail-safe, the BIG-IP system monitors various hardware components, as well as the heartbeat of various system services, and takes action if the system detects a failure.

Configuring hardware-component monitoring

You can configure the BIG-IP system to monitor the switch board component and then take some action if the BIG-IP system detects a failure.

Using the Configuration utility, you can specify the action that you want the BIG-IP system to take when the component fails. Possible actions that the BIG-IP system can take are:

- Reboot the BIG-IP system.
- Restart all system services.
- Fail over to the peer system.
- Take no action.

To configure hardware-component monitoring

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. From the Fail-safe menu, choose System.
The System Fail-safe screen opens.
3. In the System Trigger Properties area, for **Hardware Monitor**, verify that the box is checked.
4. From the **Switch Board Failure** list, select an action, or retain the default setting (**Fail Over**).
5. Click **Update**.

Configuring system-service monitoring

You can configure the BIG-IP system to monitor various system services and then take some action if the BIG-IP system detects a heartbeat failure. These services are:

- MCPD (messaging and configuration)
- TMM (traffic management)
- BIGD (health monitors)
- SOD (failover)
- BCM56XXD (switch hardware driver)

Using the Configuration utility, you can specify the action that you want the BIG-IP system to take when the heartbeat of a system service fails. Table 12.2 lists each system service, and shows the possible actions that the BIG-IP system can take in the event of a heartbeat failure.

System Service	Possible actions
MCPD	Restart the system service. Restart all system services.
TMM	Reboot the BIG-IP system. Fail over to the peer system. Fail over to the peer system and restart the service.
BIGD	Take no action.
SOD	Reboot the system. Restart all system services. Take no action.
BCM56XXD	Restart the system service. Take no action.

Table 12.2 Possible actions in response to heartbeat failure

To configure system-service monitoring

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. From the Fail-safe menu, choose System.
The System Fail-safe screen opens.
3. In the System Services area, in the Name column, click the name of the service you want to monitor.
The screen for that service opens.
4. From the **Heartbeat Failure** list, select an action, or retain the default setting.
5. Click **Finished**.
The System Fail-safe screen opens.
6. Click **Update**.

Configuring gateway fail-safe

Fail-safe features on the BIG-IP system provide network failure detection based on network traffic. One type of network failure detection is known as gateway fail-safe. *Gateway fail-safe* monitors traffic between the active BIG-IP system and a pool containing a gateway router, thereby protecting the system from a loss of an internet connection by triggering a failover when a gateway router is unreachable for a specified duration. If you want failover to occur when a gateway router is unreachable, you can configure the gateway fail-safe feature.

You can configure gateway fail-safe using the Configuration utility. Configuring gateway fail-safe means designating a load balancing pool as a gateway fail-safe pool.

When you designate a pool as a gateway fail-safe pool, you provide the following information:

- Name of the pool
- The unit ID number of the peer on which the gateway pool is configured
- The minimum number of gateway pool members that must be available to avoid the designated action
- The action that the BIG-IP system should take when the number of available gateway pool members drops below the designated threshold. Possible actions are **Reboot**, **Fail Over**, and **Restart All**. The default value is **Fail Over**.

To configure gateway fail-safe

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. From the Fail-safe menu, choose Gateway.
The Gateway Fail-safe screen opens.
3. In the upper-right corner of the screen, click **Add**.
The Add Gateway Pool screen. opens.
4. From the **Gateway Pool** list, select the name of a load balancing pool.
5. From the **Unit ID** list, select a unit ID for the pool (1 or 2).
6. In the **Threshold** box, type the number of pool members that must be available to avoid the action designated in the **Action** setting.
7. From the **Action** list, select an action (**Reboot**, **Fail Over**, or **Restart All**).
8. Click **Finished**.

Configuring VLAN fail-safe

For maximum reliability, the BIG-IP system supports failure detection on all VLANs. When you configure the fail-safe option on a VLAN, the BIG-IP system monitors network traffic going through a VLAN. If the BIG-IP system detects a loss of traffic on a VLAN and the fail-safe timeout period has elapsed, the BIG-IP system attempts to generate traffic by issuing ARP requests to nodes accessible through the VLAN. The BIG-IP system also generates an ARP request for the default route, if the default router is accessible from the VLAN. Failover is averted if the BIG-IP system is able to send and receive any traffic on the VLAN, including a response to its ARP request.

If the BIG-IP system does not receive traffic on the VLAN before the timeout period expires, it can either initiate failover and switch control to the standby unit, reboot, or restart all system services. The default action is **Restart All**.

◆ WARNING

You should configure the fail-safe option on a VLAN only after the BIG-IP system is in a stable production environment. Otherwise, routine network changes might cause failover unnecessarily.

Each interface card installed on the BIG-IP system is typically mapped to a different VLAN. Thus, when you set the fail-safe option on a particular VLAN, you need to know the interface to which the VLAN is mapped. You can use the Configuration utility to view VLAN names and their associated interfaces.

There are two ways to configure VLAN fail-safe: from the Redundancy Properties screen, or from the VLANs screen.

To configure VLAN fail-safe using the Redundancy Properties screen

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. From the Fail-safe menu, choose VLANs.
The VLAN Fail-safe screen opens.
3. In the upper-right corner of the screen, click **Add**.
The Add VLAN screen opens.
4. From the **VLAN** list, select a VLAN name.
5. In the **Timeout** box, specify the period of time during which traffic should be detected on the VLAN, after which the designated action will occur.
The default value, in seconds, is **30**.
6. From the **Action** list, select the action that the BIG-IP system should take when the timeout period expires.
Possible actions are **Reboot**, **Fail Over**, and **Restart All**.
7. Click **Finished**.

To configure VLAN fail-safe using the VLANs screen

1. On the Main tab of the navigation screen, expand **Network**, and click **VLANs**.
This opens a list of existing VLANs.
2. Click the name of an existing VLAN, or click **Create**.
3. For the **Configuration** heading, select **Advanced**.
4. In the **Fail-safe** setting, check the box.
This shows additional settings.
5. In the **Fail-safe Timeout** box, specify the period of time during which traffic should be detected on the VLAN, after which the designated action will occur.
The default value, in seconds, is **30**.
6. From the **Action** list, select the action that the BIG-IP system should take when the timeout period expires.
Possible actions are **Reboot**, **Fail Over**, and **Restart All**.
7. Click **Update** or **Finished**.

Mirroring connection information

The failover process of a redundant system ensures that a BIG-IP system is always available to process connections with no discernible interruption in service at the time of the failure event. However, failover does not preserve the specific connections on your servers at the moment of failover; connections are dropped when an active unit becomes unavailable, unless you have enabled connection mirroring.

The *connection mirroring* feature on the BIG-IP system duplicates a unit's state (that is, real-time connection and persistence information) on the peer unit. When connection mirroring is enabled, failover can be so seamless that file transfers can proceed uninterrupted and your servers can generally continue with whatever they were doing at the time of failover.

◆ **Note**

You cannot mirror Secure Sockets Layer (SSL) connections.

To enable connection mirroring for a virtual server

1. On the Main tab of the navigation pane, expand **Local Traffic**, and click **Virtual Servers**.
This opens the Virtual Servers screen, which displays a list of existing virtual servers.
2. Click the name of a virtual server.
This shows the properties for that virtual server.
3. Next to the Configuration heading, select **Advanced**.
4. Scroll down to the **Connection Mirroring** setting and check the box.
5. Click **Update**.

To enable connection mirroring for a SNAT

1. On the Main tab of the navigation pane, expand **Local Traffic**, and click **SNATs**.
This opens the SNATs screen, which displays a list of existing SNATs.
2. Click the name of a SNAT.
This displays the properties for that SNAT.
3. In the Configuration area, locate the **Stateful Failover Mirror** setting and check the box.
4. Click **Update**.

In addition to setting up connection mirroring, you can change the values of several bigdb database keys related to mirroring. For more information, see the man page for the **bigpipe db** command.

Setting a shared MAC masquerade address

For active/standby systems only, you can share the media access control (MAC) masquerade address between BIG-IP units. This feature is useful if you want to use the system in a topology with secure hubs. Sharing the MAC masquerade address between units has the following advantages:

- Increased reliability and failover speed, especially in lossy networks
- Interoperability with switches that are slow to respond to the network changes
- Interoperability with switches that are configured to ignore network changes

A BIG-IP unit uses the shared MAC address when it is the active unit in a redundant system. When the unit is in a standby state, the system uses the actual MAC address of the interface. Note that this option is valid only for active/standby mode.

◆ **Note**

For sensible operation, you must set the MAC masquerade address to be the same on both the active and standby units.

Viewing interfaces and MAC addresses

The MAC address for a VLAN is the MAC address of the first interface to be mapped to the VLAN, typically 4.1 for external and 5.1 for internal. You can view the interfaces mapped to a VLAN using the Configuration utility.

To view the interfaces mapped to a VLAN

1. On the Main tab of the navigation pane, expand **Network**, and click **VLANs**.
This opens the VLANs screen, which displays a list of existing VLANs.
2. In the Untagged Interfaces or Tagged Interfaces column, view the interfaces mapped to each VLAN.

You can also view the MAC addresses for the interfaces on the BIG-IP system.

To view the MAC address for an interface

1. On the Main tab of the navigation pane, expand **Network**, and click **Interfaces**.
This opens the Interfaces screen, which displays a list of existing interfaces.
2. In the MAC Address column, view the MAC address for each interface.

Designating a shared MAC masquerade address

The first step in designating a MAC masquerade address that the redundant units can share is to find the MAC address on both the active and standby units, and choose an address that is similar but unique. A safe technique for selecting the shared MAC address follows.

Suppose you want to set up **mac_masq** on the external interfaces. Using the Configuration utility on the active and standby units, you note that their MAC addresses are:

Active: 3.1 = 0:0:0:ac:4c:a2

Standby: 3.1 = 0:0:0:ad:4d:f3

To avoid packet collisions, you now must choose a unique MAC address. The safest way to do this is to select one of the addresses and logically **OR** the first byte with **0x40**. This makes the MAC address a locally-administered MAC address.

In this example, either **40:0:0:ac:4c:a2** or **40:0:0:ad:4d:f3** is a suitable shared MAC address to use on both BIG-IP units in the redundant system.

The shared MAC address is used only when the BIG-IP system is in active mode. When the unit is in a standby state, the original MAC address of the network card is used.

Use the following procedure to set the MAC masquerade address that is shared by both BIG-IP units in the redundant system.

To specify a shared MAC masquerade address

1. On the Main tab of the navigation pane, expand **Network**, and click **VLANs**.
This opens the VLANs screen, which displays a list of existing VLANs.
2. Click a VLAN name.
This opens the properties screen for that VLAN.

3. Locate the Configuration section and select **Advanced**.
This displays the MAC address for the first interface added to the VLAN, as well as a box for designating the MAC masquerade address.
4. Type the MAC masquerade address.
5. Click **Update**.

If you do not configure a MAC masquerade address on startup, or when transitioning from a standby state to an active state, the BIG-IP system sends gratuitous ARP requests to notify the default router and other machines on the local Ethernet segment that its MAC address has changed. See RFC 826 for more details on ARP.

Maintaining a redundant system

In the section titled *Before you begin*, on page 12-8, you learned how to use the Configuration utility to configure a redundant system. After you have configured your redundant system and put it into service, however, you can perform a number of tasks on a regular basis to keep the system up and running properly. These tasks are:

- Viewing the redundancy state of a unit and the synchronization state
- Changing the redundancy state of a unit
- Determining the ID of a unit
- Altering failover behavior
- Controlling failback
- Converting an active-active system to an active/standby system

Viewing redundancy states and synchronization status

You can globally view the current redundancy state of a unit by viewing the upper-left corner of any Configuration utility screen. You can also view synchronization status in this same portion of each screen, but only if you have configured the unit to do so. (For more information on viewing synchronization status, see *Viewing synchronization status*, on page 12-33.)

Figure 12.3 shows an example of the status information as displayed on a Configuration utility screen.



Figure 12.3 Viewing redundancy state and synchronization status

Viewing redundancy state

Each unit in a redundant system is in either an active or a standby state at any given time. The current state of a unit depends on the unit's redundancy state preference, and whether the unit is available or unavailable to accept connections. You can use the Configuration utility in two different ways to determine the current state of a unit.

One way to view the redundancy state of a BIG-IP unit is by checking the status display that appears in the upper-left corner of every Configuration utility screen. See Figure 12.3, on page 12-32 for an example of this status display.

The other way to view the redundancy state of a unit is to use the following procedure.

To view the redundancy state of a unit

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. View the value of the **Current Redundancy State** setting, either **Active** or **Standby**.

Viewing synchronization status

As you learned in *Synchronizing configuration data*, on page 12-18, it is essential that an active unit share its current configuration data with its peer to ensure that failover can occur successfully. Configuration data on a unit can change for a variety of reasons (such as when adding a virtual server or modifying a profile), so it is important that you be able to monitor synchronization status at any given time, and re-synchronize the units if necessary.

You can view either detailed synchronization information about the current unit and its peer, or you can view general synchronization status on every Configuration utility screen.

Viewing detailed synchronization status

You can view detailed configuration synchronization status, such as the internal static self IP addresses that the two units use when synchronizing data, by displaying the ConfigSync screen of the Configuration utility. This screen is available from the Redundancy Properties screen.

Table 12.3 lists and describes the status-related settings that the ConfigSync screen displays.

Status Setting	Description
ConfigSync Peer	Displays the internal static self IP address that the BIG-IP system uses to determine the peer unit for synchronization.
Status message	Displays the state of the peer, that is, active or standby. May also report the status as unknown if connectivity problems exist.
Last Change (Self)	Displays the day, date, and exact time that the configuration data of the unit you are configuring was last changed.
Last Change (Peer)	Displays the day, date, and exact time that the configuration data of the peer unit was last changed.
Last ConfigSync	Displays the day, date, and exact time that a user synchronized the configuration data between the two units.

Table 12.3 Status information on the ConfigSync screen

◆ **Tip**

*You can also display this detailed status information by enabling the auto detection feature and then clicking on the resulting synchronization status in the upper-left corner of the Configuration utility. For more information, see **Viewing general synchronization status**, following.*

Viewing general synchronization status

The BIG-IP system can display general synchronization status in the upper-left corner of every Configuration utility screen (see Figure 12.3, on page 12-32). To use this feature, you must first enable the feature that automatically detects synchronization status. For more information on how to enable this feature, see *Performing configuration synchronization*, on page 12-19.

Once you have enabled the automatic status detection feature, the status that the Configuration utility displays indicates whether or not you need to synchronize the configurations of the two units due to changes that you might have made to the configuration data. For example, if you change the redundancy preference of unit 2, you must then synchronize the configuration so that both units are aware of the change.

The synchronization status messages that the BIG-IP system can report are:

- **ConfigSync: OK** (green circular arrow)
No synchronization required
- **Sync Recommended** (orange/yellow circular arrow)
Synchronize to or from the peer

- **Sync Recommended** (gray/white circular arrow)
Manual intervention recommended
- **Unknown sync state** (gray/white circular arrow)
State of unit is unknown due to loss of peer connection

Note that the color of the icon changes depending on the synchronization status. To see an example of general synchronization status, see Figure 12.3, on page 12-32.

◆ **Tip**

*You can view more detailed synchronization information by clicking the synchronization status that is displayed in the upper-right corner of every Configuration utility screen. You can also view this detailed information by displaying the ConfigSync screen within the Configuration utility. For more information, see the previous section, **Viewing detailed synchronization status**.*

Changing the redundancy state

In addition to viewing the current state of a BIG-IP unit, you can force a BIG-IP unit to switch from an active state to a standby state. This feature is useful when you need to take an active unit out of service in order to perform standard system maintenance tasks on that unit.

For active/standby configurations, the forcing of the active unit to a standby state causes its connections to fail over to the peer unit, and the normally-idle peer unit becomes active.

For active-active configurations, you cannot switch an active unit to a standby state if that unit is currently processing connections for the peer unit due to a failover.

To change the redundancy state from active to standby

To force a unit from an active to a standby state, display the Redundancy Properties screen for the unit, and click the **Force to Standby** button.

Determining the unit ID

Sometimes, you might need to know the unit ID assigned to a BIG-IP unit. A typical reason for needing this information is if you are configuring a second BIG-IP system to be part of a redundant system and want to know the unit ID of the existing redundant-system unit.

The **Unit ID** setting on the Redundancy Properties screen indicates whether a system is designated as unit **1** or unit **2** in the redundant system configuration.

You assign a unit ID to a unit of a redundant system at the time that you initially run the Setup utility on one of the BIG-IP units. For more information, see *Before you begin*, on page 12-8.

To determine the unit ID of a BIG-IP unit

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. Locate the **Unit ID** setting and view the unit ID number.

Customizing redundant-system behavior

The bigdb™ database contains a number of bigdb keys related to redundant-system behavior. You can change the values of these keys if you want to customize the way that a redundant system operates. For more information, see Appendix B, *Configuring bigdb Database Keys*.

Controlling failback

Failback is the process of a previously-unavailable unit reclaiming its normal traffic when the unit returns to an active state. The way that failback operates on a redundant system depends on whether your system is an active/standby or an active-active configuration.

When failback occurs in either an active/standby or active-active configuration, the BIG-IP system drops any connections that had previously failed over to the peer unit, unless you configured the relevant virtual server or SNAT to mirror those connections. For information on connection mirroring, see *Mirroring connection information*, on page 12-28.

◆ Tip

*The upper-left corner of the Configuration utility screens continually reports the state of the redundant-system unit that you are currently managing (active or standby). For more information, see **Viewing redundancy state**, on page 12-32.*

Controlling failback for active/standby mode

In a typical active/standby configuration, where no preference is set for which unit is active, failback is not automatic. If unit 1 fails over to unit 2 and later becomes available again, unit 1 remains in a standby state and unit 2 continues to process the traffic, unless you specifically initiate failback or failover occurs again. In the latter case, unit 2 fails over to unit 1.

You normally initiate failback by forcing the currently-active unit into a standby state. When you initiate failback, the BIG-IP system drops the connections that failed over from the other unit, unless you have connection

mirroring enabled on the relevant virtual server or SNAT. For information on connection mirroring, see *Mirroring connection information*, on page 12-28.

As an option, you can set a state preference for a particular unit. Setting this preference automates the failback process. For example, if you specify that you prefer unit 1 to be the active unit whenever possible, then automatic failback to unit 1 occurs as soon as possible after a failover to unit 2. Automatic failback also occurs if you specify that you prefer unit 2 to be the standby unit. For information on setting a state preference see *Specifying a redundancy state preference*, on page 12-14.

To initiate failback for an active/standby system

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. At the bottom of the screen, click the **Force to Standby** button.

Controlling failback for active-active mode

Failback for an active-active configuration is never automatic, because during failback, the BIG-IP system drops the connections that failed over from the other unit, unless you have connection mirroring enabled on the relevant virtual server. For information on connection mirroring, see *Mirroring connection information*, on page 12-28.

You can, however, specifically initiate failback. In this case, failback causes the failover unit (for example, unit 2) to relinquish the processing of unit 1 traffic back to unit 1, while continuing to process its own traffic.

To initiate failback for an active-active system

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. At the bottom of the screen, click the **Fail Back** button.

◆ Note

*You can enable automatic failback by configuring the **Common.Bigip.ManFailBack** bigdb database key. For more information, see Appendix B, *Configuring bigdb Database Keys*.*

Converting an active-active system to an active/standby system

Returning to active/standby mode from active-active mode is relatively simple in that only a few things need be changed. Use the Configuration utility to do this conversion, being sure to make these same changes on each unit of the redundant pair.

To change the redundancy mode of a redundant system

1. On the Main tab of the navigation pane, expand **System**, and click **High Availability**.
The Redundancy Properties screen opens.
2. From the **Redundancy Mode** list, change the mode to **Active/Standby**.
3. On the menu bar, click ConfigSync.
The ConfigSync screen opens.
4. For the **Synchronize** setting, click **Synchronize TO Peer**.
5. Click **Update**.
6. On the menu bar, click Redundancy.
The Redundancy Properties screen opens.
7. From the **Redundancy State Preference** list, change the preferred state of the unit to **Active** or **Standby**. For more information, see *Specifying a redundancy state preference*, on page 12-14.
8. Click **Update**.

Once the system is running in active/standby mode, it is the active unit that manages connections for all local traffic management objects (such as virtual servers and SNATs), regardless of the unit associated with those objects. Therefore, it is not necessary to re-associate virtual servers, SNATS, or NATs with a specific unit when you transition from active-active mode to active/standby mode.



A

Troubleshooting SNMP Traps

- Understanding F5-specific traps

Understanding F5-specific traps

The alert system and the SNMP agent on the BIG-IP system can send a number of notifications about the current operation of the BIG-IP system. The definitions of these notifications are known as *traps*, and are contained in the MIB file **F5-BIGIP-COMMON-MIB.txt**.

The types of notifications that the BIG-IP system can send to an SNMP manager are:

- General traps
- Hardware-related traps
- License-related traps
- Traffic management operating system (TMOS)-related traps
- Authentication-related traps
- Denial of Service (DoS)-related traps
- Network-related traps
- Logging-related traps

The remainder of this appendix contains tables that list the trap names contained in the MIB file **F5-BIGIP-COMMON-MIB.txt**, their descriptions, and the recommended action to take if you receive one of these notifications on your SNMP manager system.

General traps

The general notifications that you might receive on an SNMP manager system are listed and described in Table A.1.

Trap Name	Description	Recommended Action
bigipAgentStart	The SNMP agent on the BIG-IP system has been started.	For your information only. No action required.
bigipAgentShutdown	The SNMP agent on the BIG-IP system has been shut down.	
bigipAgent Restart	The SNMP agent on the BIG-IP system has been restarted.	
bigipConfigLoaded	The BIG-IP system configuration data was loaded.	
bigipDiskPartitionWarn	Free space on the disk partition is limited.	Increase the available disk space.
bigipDiskPartitionGrowth	The disk partition use exceeds the growth limit.	

Table A.1 General traps and recommended actions

Hardware-related traps

The hardware-related notifications that you might receive on an SNMP manager system are listed and described in Table A.2.

Trap Name	Description	Recommended Action
bigipCpuTempHigh	The CPU temperature is too high.	Contact F5 Networks technical support.
bigipCpuFanSpeedLow	The CPU fan speed is too low.	
bigipCpuFanSpeedBad	The CPU fan is not receiving a signal.	
bigipChassisTempHigh	The temperature of the chassis is too high.	
bigipChassisFanBad	The chassis fan is not operating properly.	
bigipChassisPowerSupplyBad	The chassis power supply is not functioning properly.	

Table A.2 Hardware-related traps and recommended actions

License-related traps

The notifications related to licensing that you might receive on an SNMP manager system are listed and described in Table A.3.

Trap Name	Description	Recommended Action
bigipLicenseFailed	Validation of a BIG-IP system license has failed, or the dossier has errors. Occurs only when first licensing the system or adding a module key (such as HTTP compression) to an existing system.	If using automatic licensing, verify connectivity to the outside world, fix the dossier if needed, and try again.
bigipLicenseExpired	The BIG-IP license has expired.	Call F5 Networks technical support.

Table A.3 License traps and recommended actions

TMOS-related traps

The TMOS-related notifications that you might receive on an SNMP manager system are listed and described in Table A.4.

Trap Name	Description	Recommended Action
bigipServiceDown	A BIG-IP system health monitor has detected a service on a node to be stopped and has therefore marked the node as down .	Restart the service on the node.
bigipServiceUp	A BIG-IP system health monitor has detected a service on a node to be running and has therefore marked the node as up .	For your information only. No action required/
bigipNodeDown	A BIG-IP system health monitor has marked a node as down .	Check the node and the cable connection.
bigipNodeUp	A BIG-IP system health monitor has marked a node as up .	For your information only. No action required.
bigipStandby	The BIG-IP system has switched to standby mode.	Review the log files in the /var/log directory and then search for core files in the /var/core directory. If you find a core file, or find text similar to fault at location xxx stack trace: , contact F5 Networks technical support.
bigipActive	The BIG-IP system has switched to active mode.	For your information only. No action required.
bigipActiveActive	The BIG-IP system is in active-active mode.	
bigipFeatureFailed	A high-availability feature has failed.	View high-availability processes and their current status using the BIG-IP system command bigpipe ha table .
bigipFeatureOnline	A high-availability feature is responding.	For your information only. No action required.
bigipPacketRejected	The BIG-IP system has rejected some packets.	Check the detailed message within this trap and act accordingly.

Table A.4 TMOS-related traps and recommended actions

Authentication-related traps

The notifications related to authentication that you might receive on an SNMP manager system are listed and described in Table A.5.

Trap Name	Description	Recommended Action
bigipTamdAlert	More than 60 authentication attempts have failed within one second, for a given virtual server.	Investigate for a possible intruder.
bigipAuthFailed	A login attempt failed.	Check the user name and password.

Table A.5 Authentication traps and recommended actions

DoS-related traps

The notifications related to denial of service that you might receive on an SNMP manager system are listed and described in Table A.6.

Trap Name	Description	Recommended Action
bigipAggrReaperStateChange	The state of the aggressive reaper has changed, indicating that the BIG-IP system is moving to a distress mode.	We recommend that you use our default Denial of Service settings. However, you can add rate filters to survive the attack.

Table A.6 Denial of Service traps and recommended actions

Network-related traps

The network-related notifications that you might receive on an SNMP manager system are listed and described in Table A.7.

Trap Name	Description	Recommended Action
bigipARPConflict	The BIG-IP system has detected an ARP advertisement for any of its own ARP-enabled addresses. This can occur for a virtual server address or a self IP address.	Check IP addresses and routes.
bigipNetLinkDown	An interface link is down.	Check the cable connection.

Table A.7 Network traps and recommended actions

Logging-related traps

The notifications related to logging that you might receive on an SNMP manager system are listed and described in Table A.8.

Trap Name	Description	Recommended Action
bigipLogEmerg	The BIG-IP system is unusable. This notification occurs when the system logs a message with the log level LOG_EMERG .	Check the detailed message within this trap and within the /var/log files to determine which process has the emergency. Then act accordingly.
bigipLogAlert	You must take action immediately for the BIG-IP system to function properly. This notification occurs when the system logs a message with the log level LOG_ALERT .	Check the detailed message within this trap and within the /var/log files to determine which process has the alert situation. Then act accordingly.
bigipLogCrit	The BIG-IP system is in a critical condition. This notification occurs when the system logs a message with the log level LOG_CRIT .	Check the detailed message within this trap and within the /var/log files to determine which process has the critical situation. Then act accordingly.
bigipLogErr	The BIG-IP system has some error conditions. This notification occurs when the system logs a message with the log level LOG_ERR .	Check the detailed message within this trap and within the /var/log files to determine which processes have the error conditions. Then act accordingly.
bigipLogWarning	The BIG-IP system is experiencing some warning conditions. This notification occurs when the system logs a message with the log level LOG_WARNING .	Check the detailed message within this trap and within the /var/log files to determine which processes have the warning conditions. Then act accordingly.

Table A.8 Logging-related traps and recommended actions

◆ Note

You can also view logging information using the Logs screen of the Configuration utility.



B

Configuring bigdb Database Keys

- Introducing the bigdb database
- Summarizing bigdb keys for redundant system administration
- Summarizing bigdb keys for user account administration

Introducing the bigdb database

Every BIG-IP system includes a bigdb™ database. The bigdb database holds a set of *bigdb configuration keys*, which define the behavior of various aspects of the BIG-IP system. For example, the bigdb key **Failover.Active Mode**, when set to **enable**, causes a redundant system to operate in active-active mode, instead of the default active/standby mode.

You can change the value of a bigdb key in two ways:

- ◆ **The Configuration utility**

When you use the Configuration utility to configure various BIG-IP features, you are actually resetting bigdb key values. In this case, the bigdb keys are invisible to users.

- ◆ **The bigpipe db command**

You can reset bigdb key values directly using the **bigpipe db** command. This command is useful if you prefer not to use the Configuration utility to configure a BIG-IP feature, or if configuration of a particular aspect of BIG-IP system behavior is not available through the Configuration utility. The syntax for displaying and setting bigdb keys is:

```
bigpipe db all list
bigpipe db <key name> <value>
```

- ◆ **Tip**

*For more information on using the **bigpipe db** command, see the online man page for the command.*

Some of the bigdb database keys for system management that you might want to configure pertain specifically to redundant systems and user accounts.

Summarizing bigdb keys for redundant system administration

There are several bigdb keys that you can use to configure and manage a redundant system. These keys pertain to the following redundant-system features:

- Failover
- Connection mirroring
- Configuration synchronization
- System fail-safe

Using failover keys

The bigdb keys that you can configure for failover are shown in Table B.1, on page B-3. These keys are listed in alphabetical order.

Key Name	Default Value	Description
Failover.ActiveMode	disable	Enables or disables active-active mode. Use active-active mode if set to 1 . By default, this is 0 (off) and active/standby mode is used. Possible values are enable and disable .
Failover.DbgFile	/var/log/sodlog	Specifies the file into which the sod service logs the failover debug information.
Failover.FailbackDelay	60	For an active-active system, when the failed unit becomes active again, specifies the number of seconds that you want the system to wait before failback occurs.
Failover.FailedStandbyActive	disable	Controls whether a standby unit with a failover condition becomes active when the peer unit fails. Possible values are enable and disable .
Failover.ForceActive	disable	Specifies that the failover daemon should always attempt to become the active unit. Possible values are enable and disable .
Failover.ForceStandby	disable	Specifies that the failover daemon should switch to a standby state whenever the current unit senses that its peer is alive. Possible values are enable and disable .
Failover.IsRedundant	false	Defines whether the BIG-IP system is a unit of a redundant pair. Possible values are true and false .
Failover.ManFailBack	disable	If using active-active mode, specifies that the system should wait until the surviving unit receives a command before surrendering resources to a rebooted machine. Possible values are enable and disable .
Failover.MemoryRestartPercent	97	Defines the amount of memory usage that causes the BIG-IP system to reboot.
Failover.Network	0	Specifies whether the system should use the network as a backup to, or instead of, the hard-wired connection for failover. Possible values are 0 (off) and 1 (on).
Failover.PrintPeerState	disable	Specifies that the failover daemon (/sbin/sod) should write the state of its connection (hard-wired or network) to its peer. The system writes this information to the failover daemon's debug log file. Possible values are enable and disable .
Failover.Standby.LinkDownTime	0	Defines the amount of time in seconds that the system's interfaces are down before switching to a standby state.
Failover.UnitId	1	Specifies the ID of the unit. Each BIG-IP system must have a unique unit ID of 1 or 2 in the event that network communication is not possible with its peer.

Table B.1 bigdb database keys pertaining to failover

Key Name	Default Value	Description
Failover.UseTty00	disable	Specifies that the failover daemon should use /dev/tty00 for hard-wired failover. Possible values are enable and disable .
Failover.UseTty01	disable	Specifies that the failover daemon should use /dev/tty01 for hard-wired failover. Possible values are enable and disable .

Table B.1 *bigdb database keys pertaining to failover*

Using connection mirroring keys

The bigdb keys that you can configure for connection mirroring are shown in Table B.2. These keys are listed in alphabetical order.

Key Name	Default Value	Description
StateMirror.Ipaddr	No default value	Specifies the unit's primary static self IP address that its peer uses to mirror connections.
StateMirror.PeerIpaddr	No default value	Specifies the peer unit's primary static self IP address that a unit uses to mirror connections.
StateMirror.PeerListenPort	1028	Defines the port on which the BIG-IP system listens for connections from the active unit.
StateMirror.Secondary.Ipaddr	No default value	Specifies the unit's secondary static self IP address that its peer uses to mirror connections.
StateMirror.SecondaryPeerIpaddr	No default value	Specifies the peer unit's secondary static self IP address that a unit uses to mirror connections.
StateMirror.State	enable	Defines whether connection mirroring is enabled or disabled for a redundant system. Possible values are enable and disable .

Table B.2 *bigdb database keys pertaining to connection mirroring*

Using configuration synchronization keys

The bigdb keys that you can configure for synchronizing configuration data are shown in Table B.3. These keys are listed in alphabetical order.

Key Name	Default Value	Description
Configsync.Autodetect	enable	Defines whether the Configuration utility should automatically detect configuration status and display it on all Configuration utility screens. Possible values are enable and disable .
Configsync.LocalConfigTime	0	Specifies the most recent date and time that the configuration of the current unit changed.
Configsync.LocalSyncedTime	0	Specifies the date and time that the configuration of this unit was synchronized with the peer unit.
Configsync.password	No default value	Defines the password of the user account that has permission to synchronize configuration data.
Configsync.PeerConfigTime	0	Specifies the most recent date and time that the configuration of the peer unit changed.
Configsync.PeerState	unknown	Defines whether the peer's synchronization state is known. Possible values are known and unknown .
Configsync.PeerUpdatedTime	0	Specifies the date and time that this unit successfully informed its peer of a configuration change on this unit.
Configsync.State	-1	Specifies the configuration state of this box. Possible values are: -1 - Uninitiated or disabled config state. 0 - Synchronized. 1 - Configuration on current unit was modified. Recommend configuration synchronization to peer unit. 2 - Configuration on peer unit was modified. Recommend configuration synchronization from peer unit. 3 - Configuration modified on both units. Manual intervention required.
Configsync.username	admin	Defines the user account that has permission to synchronize configuration data.

Table B.3 bigdb database keys pertaining to configuration synchronization

Using system fail-safe keys

The bigdb keys that you can configure for system fail-safe are shown in Table B.4. These keys are listed in alphabetical order.

Key Name	Default Value	Description
Switchboard.Failsafe	enable	Enables or disables fail-safe when the switch board fails. Possible values are enable and disable .
Switchboard.Failsafe.Action	failover	Specifies the action that the system takes when the switch board fails. Possible values are failover , reboot , and restart_all .

Table B.4 bigdb database keys pertaining to system fail-safe

Summarizing bigdb keys for user account administration

You can configure a set of bigdb keys to manage administrative user accounts for a BIG-IP system. These keys and their descriptions appear in Table B.5, and are listed in alphabetical order.

Key Name	Default Value	Description
User.AcceptedEULA	none	Specifies fields that the Setup utility populates. Possible values are none , internal , non-production , and production .
Users.Default.Role	127	Specifies a numeric value for the default role for remote user accounts.
Users.LocalOnly	root,admin	Specifies those user accounts that must reside locally on the BIG-IP system and therefore cannot reside on a remote authentication server.
Users.Name.admin	0	Specifies a numeric value for the admin account.
Users.Name.support	0	Specifies a numeric value for the support account.
Users.Name.[user name]	127	Specifies a numeric value for any user account that is not root , admin , or support .

Table B.5 bigdb database keys pertaining to user accounts



Glossary

active unit

In a redundant system, the active unit is the system that currently load balances connections. If the active unit in the redundant system fails, the standby unit assumes control and begins to load balance connections. See also *redundant system*.

ARP cache

The ARP (address resolution protocol) cache is the mechanism that a device on a network uses to determine the MAC address of another device, when only the IP address of that other device is known.

authentication

Authentication is the process of verifying a user's identity when the user is attempting to log on to a system.

authentication iRule

An authentication iRule is a system-supplied or user-created iRule that is necessary for implementing a PAM authentication module on the LTM system. See also *iRule*, *PAM (Pluggable Authentication Module)*.

authentication module

An authentication module is a PAM module that you create to perform authentication or authorization of client traffic. See also *PAM (Pluggable Authentication Module)*.

authentication profile

An authentication profile is a configuration tool that you use to implement a PAM authentication module. Types of authentication modules that you can implement with an authentication profile are: LDAP, RADIUS, TACACS+, SSL Client Certificate LDAP, and OCSP. See also *PAM (Pluggable Authentication Module)*.

authorization

Authorization is the process of identifying the level of access that a logged-on user has been granted to system resources.

bigtop

The **bigtop** utility is a statistical monitoring utility that ships on the BIG-IP system. This utility provides real-time statistical information.

BIND (Berkeley Internet Name Domain)

BIND is the most common implementation of the Domain Name System (DNS). BIND provides a system for matching domain names to IP addresses. For more information, refer to <http://www.isc.org/products/BIND>.

BPDU (bridge protocol data unit)

A BPDU is a special packet that a spanning tree protocol sends between layer 2 devices to determine redundant paths, and provide loop resolution. See also *STP (Spanning Tree Protocol)*, *RSTP (Rapid Spanning Tree Protocol)*, and *MSTP (Multiple Spanning Tree Protocol)*.

bridge

A bridge is a layer 2 networking device that defines a collision domain.

bursting

Bursting is an aspect of rate shaping and occurs when the rate of traffic flow exceeds the base rate defined.

certificate

A certificate is an online credential signed by a trusted certificate authority and used for SSL network traffic as a method of authentication.

certificate authority (CA)

A certificate authority is an external, trusted organization that issues a signed digital certificate to a requesting computer system for use as a credential to obtain authentication for SSL network traffic.

certificate revocation list (CRL)

A certificate revocation list is a list that an authenticating system checks to see if the SSL certificate that the requesting system presents for authentication has been revoked.

certificate verification

Certificate verification is the part of an SSL handshake that verifies that a client's SSL credentials have been signed by a trusted certificate authority.

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.

chunking

See *HTTP chunking*.

cipher

A cipher is an encryption/decryption algorithm that computer systems use when transmitting data using the SSL protocol.

client-side SSL profile

A client-side SSL profile is an SSL profile that controls the behavior of SSL traffic going from a client system to the LTM system.

clone pool

This feature causes a pool to replicate all traffic coming into it and send that traffic to a duplicate pool.

configuration object

A configuration object is a user-created object that the LTM system uses to implement a PAM authentication module. There is one type of configuration object for each type of authentication module that you create. See also *PAM (Pluggable Authentication Module)*.

configuration synchronization

Configuration synchronization is the task of duplicating a BIG-IP system's configuration data onto its peer unit in a redundant system.

Configuration utility

The Configuration utility is the browser-based application that you use to configure the LTM system.

connection persistence

Connection persistence is an optimization technique whereby a network connection is intentionally kept open for the purpose of reducing handshaking.

connection pooling

Connection pooling is an optimization feature that pools server-side connections for re-use by other client requests. Connection pooling reduces the number of new connections that must be opened for server-side client requests.

content switching

Content switching is the ability to select a pool based on data contained within a packet.

cookie persistence

Cookie persistence is a mode of persistence where the LTM system stores persistent connection information in a cookie.

custom profile

A custom profile is a profile that you create. A custom profile can inherit its default settings from a parent profile that you specify. See also *parent profile*.

default profile

A default profile is a profile that the LTM system supplies with default setting values. You can use a default profile as is, or you can modify it. You can also specify it as a parent profile when you create a custom profile. You cannot create or delete a default profile. See also *profile*, *custom profile*.

default route

The default route is a route that you configure on a network device. For a router device, you configure the default route in the routing table. The device uses this route to forward packets when the table contains no other entry that more closely matches the destination address. You can also configure a default route on a destination host.

default VLAN

The LTM system is configured with two default VLANs, one for each interface. One default VLAN is named **internal** and one is named **external**. See also *VLAN (virtual local area network)*.

default wildcard virtual server

A default wildcard virtual server has an IP address and port number of **0.0.0.0:0**, or ***:*** or **"any":"any"**. This virtual server accepts all traffic that does not match any other virtual server defined in the configuration.

destination address affinity persistence

Also known as sticky persistence, destination address affinity persistence supports TCP and UDP protocols, and directs session requests to the same server based solely on the destination IP address of a packet.

domain name

A domain name is a unique name that is associated with one or more IP addresses. Domain names are used in URLs to identify particular Web pages. For example, in the URL **http://www.siterequest.com/index.html**, the domain name is **siterequest.com**.

Dynamic Ratio load balancing method

Dynamic Ratio mode is like Ratio mode (see *Ratio method*), except that ratio weights are based on continuous monitoring of the servers and are therefore continually changing. Dynamic Ratio load balancing can be implemented on RealNetworks® RealServer platforms, on Microsoft® Windows® platforms equipped with Windows Management Instrumentation (WMI), or on a server equipped with either the UC Davis SNMP agent or Windows 2000 Server SNMP agent.

dynamic route

A dynamic route is a route that an advanced routing protocol such as RIP adds dynamically to a routing table. See also *static route*.

EAV (Extended Application Verification)

EAV is a health check that verifies an application on a node by running that application remotely. EAV health check is only one of the three types of health checks available on an LTM system. See also *health check*, *health monitor*, and *external monitor*.

ECV (Extended Content Verification)

ECV is a health check that allows you to determine if a node is **up** or **down** based on whether the node returns specific content. ECV health check is only one of the three types of health checks available on an LTM system. See also *health check*.

external authentication

External authentication refers to the process of using a remote server to store data for the purpose of authenticating users or applications attempting to access the LTM system.

external monitor

An external monitor is a user-supplied health monitor. See also *health check*, *health monitor*.

external VLAN

The external VLAN is a default VLAN on the BIG-IP system. 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.

failback

Failback is the process whereby an active unit relinquishes processing to a previously-failed unit that is now available.

failover

Failover is the process whereby a standby unit in a redundant system takes over when a software failure or a hardware failure is detected on the active unit.

failover cable

The failover cable directly connects the two units together in a redundant system.

Fastest method

Fastest mode is a load balancing method that passes a new connection based on the fastest response of all currently active nodes.

FDDI (Fiber Distributed Data Interface)

FDDI is a multi-mode protocol used for transmitting data on optical-fiber cables at speeds up to 100 Mbps.

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.

forwarding virtual server

A forwarding virtual server is a virtual server that has no pool members to load balance. The virtual server simply forwards the packet directly to the destination IP address specified in the client request. See also *virtual server*.

hash persistence

Hash persistence allows you to create a persistence hash based on an existing iRule.

health check

A health check is an LTM system feature that determines whether a node is **up** or **down**. Health checks are implemented through health monitors. See also *health monitor*, *ECV*, *EAV*, and *external 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**. Different monitors exist for checking different services. See also *health check*, *EAV*, *ECV*, and *external monitor*.

host virtual server

A host virtual server is a virtual server that represents a specific site, such as an Internet web site or an FTP site, and it load balances traffic targeted to content servers that are members of a pool.

HTTP chunking

HTTP chunking refers to the HTTP/ 1.1 feature known as chunked encoding, which allows HTTP messages to be broken up into several parts. Chunking is most often used by servers when sending responses.

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**.

HTTP transformation

When the LTM system performs an HTTP transformation, the system manipulates the **Connection** header of a server-side HTTP request, to ensure that the connection stays open. See also *connection persistence*.

ICMP (Internet Control Message Protocol)

ICMP is an Internet communications protocol used to determine information about routes to destination addresses.

i-mode

i-mode® is a service created by NTT DoCoMo, Inc., that allows mobile phone users access to the Internet.

interface

A physical port on a BIG-IP system is called an interface.

internal VLAN

The internal VLAN is a default VLAN on the BIG-IP system. 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.

IPSEC

IPSEC (Internet Security Protocol) is a communications protocol that provides security for the network layer of the Internet without imposing requirements on applications running above it.

iRule

An iRule is a user-written script that controls the behavior of a connection passing through the LTM system. iRules™ are an F5 Networks feature and are frequently used to direct certain connections to a non-default load balancing pool. However, iRules can perform other tasks, such as implementing secure network address translation and enabling session persistence.

iSNAT (intelligent SNAT)

An iSNAT is the mapping of one or more original client IP addresses to a translation address from within an iRule. Before writing an iRule to create an iSNAT, you must create a SNAT pool. See also *SNAT pool*.

JAR file

A JAR file is a file in Java™ Archive (JAR) file format that enables you to bundle multiple files into a single archive file. Typically, a JAR file contains the class files and auxiliary resources associated with applets and applications.

JDBC

JDBC is a Java™ technology. It is an application programming interface that provides database management system (DBMS) connectivity across a wide range of SQL databases, as well as access to other tabular data sources, such as spreadsheets or flat files.

Kilobytes/Second mode

The Kilobytes/Second mode is a dynamic load balancing mode that distributes connections based on which available server currently processes the fewest kilobytes per second.

LACP (Link Aggregation Control Protocol)

LACP is an IEEE-defined protocol that provides link aggregation within a trunk, by sending control packets to continually assess link status. LACP provides link failover if a member link becomes unavailable.

last hop

A last hop is the final hop a connection takes to get to the BIG-IP system. You can allow the BIG-IP system to determine the last hop automatically to send packets back to the device from which they originated. You can also specify the last hop manually by making it a member of a last hop pool.

layer 2 forwarding table

A layer 2 forwarding table correlates MAC addresses of network devices to the BIG-IP system interfaces through which those devices are accessible. On a BIG-IP system, each VLAN has its own layer 2 forwarding table.

LDAP (Lightweight Directory Access Protocol)

LDAP is an Internet protocol that email programs use to look up contact information from a server.

LDAP authentication module

An LDAP authentication module is a user-created module that you implement on an LTM system to authenticate client traffic using a remote LDAP server.

LDAP client certificate SSL authentication module

An LDAP client certificate SSL authentication module is a user-created module that you implement on an LTM system to authorize client traffic using SSL client credentials and a remote LDAP server.

Least Connections method

Least Connections mode is a dynamic load balancing method that bases connection distribution on which server currently manages the fewest open connections.

link aggregation

Link aggregation is the process of combining multiple links in order to function as though it were a single link with higher bandwidth. Link aggregation occurs when you create a trunk. See also *trunk* and *LACP (Link Aggregation Control Protocol)*.

load balancing method

A particular method of determining how to distribute connections across a load balancing pool.

load balancing pool

See *pool*.

load balancing virtual server

A load balancing virtual server is a virtual server that directs client traffic to a load balancing pool. This is the most basic type of virtual server. See also *virtual server*.

local traffic management (LTM)

Local traffic management (LTM) is the process of managing network traffic that comes into or goes out of a local area network (LAN), including an intranet.

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.

member

Member is a reference to a node when it is included in a particular load balancing pool. Pools typically include multiple member nodes.

minimum active members

The minimum active members is the number of members that must be active in a priority group in order for the LTM system to send its requests to that group. If the number of active members falls below this number, requests are sent to the next highest priority group (the priority group with the next lowest priority number).

monitor

The LTM system 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 association

A monitor association is an association that a user makes between a health or performance monitor and a pool, pool member, or node.

monitor instance

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

monitor template

A monitor template is an internal mechanism that the LTM system uses to provide default values for a custom monitor when no pre-configured monitor exists.

MSRDP persistence

MSRDP persistence tracks sessions between clients and servers running the Microsoft® Remote Desktop Protocol (RDP) service.

MSTP (Multiple Spanning Tree Protocol)

Defined by IEEE, MSTP is an enhanced spanning tree protocol. Unlike STP and RSTP, MSTP is VLAN-aware and therefore incorporates the concept of MSTP regions. See also *STP (Spanning Tree Protocol)* and *RSTP (Rapid Spanning Tree Protocol)*.

MSTP region

An **MSTP region** is a group of layer 2 devices that have identical values for certain configuration settings. When devices constitute a region, the spanning tree algorithm takes VLANs into account when blocking and unblocking redundant paths.

name resolution

Name resolution is the process by which a name server matches a domain name request to an IP address, and sends the information to the client requesting the resolution.

NAT (Network Address Translation)

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

network virtual server

A network virtual server is a virtual server whose IP address has no bits set in the host portion of the IP address (that is, the host portion of its IP address is **0**). There are two kinds of network virtual servers: those that direct client traffic based on a range of destination IP addresses, and those that direct client traffic based on specific destination IP addresses that the LTM system does not recognize.

node

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 LTM system uses to verify the status of multiple nodes. When the LTM system uses a node alias to check node status, it pings the node alias. If the LTM system receives a response to the ping, it marks all nodes associated with the node alias as **up**. If the LTM system does not receive a response to the ping, 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 LTM system uses the node ping and health check features to determine node status.

Observed method

Observed mode is a dynamic load balancing method that bases connection distribution on a combination of two factors: the server that currently hosts the fewest connections and also has the fastest response time.

OCSP (Online Certificate Status Protocol)

OCSP is a protocol that authenticating systems can use to check on the revocation status of digitally-signed SSL certificates. The use of OCSP is an alternative to the use of a certificate revocation list (CRL). See also *certificate revocation list (CRL)*.

OCSP authentication module

An OCSP authentication module is a user-created module that you implement on an LTM system to authenticate client traffic using a remote OCSP responder. The purpose of an OCSP authentication module is to check on the revocation status of a client SSL certificate.

OCSP responder

An OCSP responder is an external server used for communicating SSL certificate revocation status to an authentication server such as the LTM system.

OCSP responder object

A responder object is a software application on the LTM system that communicates with an OCSP responder, for the purpose of checking revocation status of a client or server SSL certificate.

OneConnect™

The F5 Networks OneConnect™ feature optimizes the use of network connections by keeping server-side connections open and pooling them for re-use.

packet rate

The packet rate is the number of data packets per second processed by a server.

PAM (Pluggable Authentication Module)

PAM is a software module that a server application uses to authenticate client traffic. The modular design of PAM allows an organization to add, replace, or remove that authentication mechanism from a server application with minimal impact to that application. An example of PAM is an application that uses a remote Lightweight Directory Access Protocol (LDAP) server to authenticate client traffic. See also *LDAP (Lightweight Directory Access Protocol)*.

parent profile

A parent profile is a profile that can propagate its values to another profile. A parent profile can be either a default profile or a custom profile. See also *profile*.

performance monitor

A performance monitor gathers statistics and checks the state of a target device.

persistence

See *connection persistence* or *session persistence*.

persistence profile

A persistence profile is a configuration tool for implementing a specific type of session persistence. An example of a persistence profile type is a cookie persistence profile.

pipelining

Pipelining is a feature of HTTP/1.0 that allows clients to make requests even when prior requests have not yet received a response from the server.

pool

A pool is composed of a group of network devices (called members). The LTM system load balances requests to the nodes within a pool based on the load balancing method and persistence method you choose when you create the pool or edit its properties.

pool member

A pool member is a server that is a member of a load balancing pool.

port

A port can be 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.

port mirroring

Port mirroring is a feature that allows you to copy traffic from any port or set of ports to a single, separate port where a sniffing device is attached.

port-specific wildcard virtual server

A port-specific wildcard virtual server is a wildcard virtual server that uses a port number other than 0. See *wildcard virtual server*.

pre-configured monitor

A pre-configured monitor is a system-supplied health or performance monitor. You can use a pre-configured monitor as is, but you cannot modify or delete one. See also *monitor*.

Predictive method

Predictive mode is a dynamic load balancing method that bases connection distribution on a combination of two factors: the server that currently hosts the fewest connections, and also has the fastest response time. Predictive method also ranks server performance over time, and passes connections to servers which exhibit an improvement in performance rather than a decline.

profile

A profile is a configuration tool containing settings for defining the behavior of network traffic. The LTM system contains profiles for managing FastL4, HTTP, TCP, FTP, SSL, and RTSP traffic, as well as for implementing persistence and application authentication.

profile setting

A profile setting is a configuration attribute within a profile that has a value associated with it. You can configure a profile setting to customize the way that the LTM system manages a type of traffic.

profile type

A profile type is a category of profile that you use for a specific purpose. An example of a profile type is an HTTP profile, which you configure to manage HTTP network traffic.

protocol profile

A protocol profile is a profile that you create for controlling the behavior of FastL4, TCP, UDP, OneConnect, and RTSP traffic.

Quality of Service (QoS) level

The Quality of Service (QoS) level is a means by which network equipment can identify and treat traffic differently based on an identifier. Essentially, the QoS level specified in a packet enforces a throughput policy for that packet.

RADIUS (Remote Authentication Dial-in User Service)

RADIUS is a service that performs remote user authentication and accounting. Its primary use is for Internet Service Providers, though it can also be used on any network that needs a centralized authentication and/or accounting service for its workstations.

RADIUS authentication module

A RADIUS authentication module is a user-created module that you implement on an LTM system to authenticate client traffic using a remote RADIUS server.

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 shaping*.

rate shaping

Rate shaping is a type of extended IP filter. Rate shaping uses the same IP filter method but applies a rate class, which determines the volume of network traffic allowed. See also *rate class*.

ratio

A ratio is a parameter that assigns a weight to a virtual server for load balancing purposes.

Ratio method

The Ratio load balancing method distributes connections across an array of virtual servers in proportion to the ratio weights assigned to each individual virtual server.

receive expression

A receive expression is the text string that the LTM system looks for in the web page returned by a web server during an extended content verification (ECV) health check.

redundant system

Redundant system refers to a pair of units that are configured for fail-over. In a redundant system, there are two units, one running as the active unit and one running as the standby unit. If the active unit fails, the standby unit takes over and manages connection requests.

reference link

A reference link is the lowest-numbered interface in a trunk and is used for link aggregation.

remote administrative IP address

A remote administrative IP address is an IP address from which a BIG-IP system allows shell connections, such as Telnet or SSH.

responder object

See *OCSP responder object*.

RFC 1918 addresses

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

Round Robin mode

Round Robin mode is a static load balancing mode that bases connection distribution on a set server order. Round Robin mode sends a connection request to the next available server in the order.

router

A router is a layer 3 networking device. If no VLANs are defined on the network, a router defines a broadcast domain.

RSTP (Rapid Spanning Tree Protocol)

Defined by IEEE, RSTP is an enhanced version of STP (Spanning Tree Protocol). RSTP provides faster spanning tree performance compared to STP. See also *STP (Spanning Tree Protocol)* and *MSTP (Multiple Spanning Tree Protocol)*.

RTSP

RTSP (Real-Time Streaming Protocol) establishes and controls one or more time-synchronized streams of continuous media such as audio or video.

self IP address

Self IP addresses are the IP addresses owned by the BIG-IP system that you use to access devices in VLANs. You assign self IP addresses to VLANs.

send string

A send string is the request that the LTM system sends to the web server during an extended content verification (ECV) health check.

server-side SSL profile

A server-side SSL profile is an SSL profile that controls SSL traffic going between an LTM system and a destination server system.

service

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

services profile

A services profile is a configuration tool on the LTM system for managing either HTTP or FTP network traffic.

session persistence

A series of related connections received from the same client, having the same session ID. When persistence is enabled, an LTM system sends all connections having the same session ID to the same node, instead of load balancing the connections. Session persistence is not to be confused with *connection persistence*.

Setup utility

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

simple persistence

See *source address affinity persistence*.

SIP persistence

SIP persistence is a type of persistence used for servers that receive Session Initiation Protocol (SIP) messages sent through UDP. SIP is a protocol that enables real-time messaging, voice, data, and video.

SNAT (Secure Network Address Translation)

A SNAT is a feature you can configure on the LTM system. 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. See also *standard SNAT* and *iSNAT (intelligent SNAT)*.

SNAT pool

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

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.

source address affinity persistence

Also known as simple persistence, source address affinity persistence supports TCP and UDP protocols, and directs session requests to the same server based solely on the source IP address of a packet.

source processing

Source processing means that the interface rewrites the source of an incoming packet.

spanning tree

A spanning tree is a logical tree structure of layer 2 devices on a network, created by a spanning tree protocol algorithm and used for resolving network loops.

spanning tree instance

A spanning tree instance is a specific, named spanning tree that a spanning tree protocol creates. See also *spanning tree protocols*.

spanning tree protocols

Spanning tree protocols are the IEEE-specified set of protocols known as Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). The BIG-IP system includes support for all of these protocols. See also *STP (Spanning Tree Protocol)*, *RSTP (Rapid Spanning Tree Protocol)*, and *MSTP (Multiple Spanning Tree Protocol)*.

SSH

SSH is a protocol for secure remote login and other secure network services over a non-secure network.

SSL (Secure Sockets Layer)

SSL is a network communications protocol that uses public-key technology as a way to transmit data in a secure manner.

SSL persistence

SSL persistence is a type of persistence that tracks non-terminated SSL sessions, using the SSL session ID.

SSL profile

An SSL profile is a configuration tool that you use to terminate and initiate SSL connections from clients and servers.

standard SNAT

A standard SNAT is a SNAT that you implement by using the SNAT screens of the Configuration utility. See also *SNAT (Secure Network Address Translation)* and *iSNAT (intelligent SNAT)*.

standby unit

A standby unit in a redundant system is a unit that is always prepared to become the active unit if the active unit fails.

state mirroring

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

static route

A static route is a route that you must explicitly configure on a layer 3 device in its routing table. See also *dynamic route*.

static self IP address

A static self IP address is a self IP address that is not shared between two units of a redundant system.

sticky persistence

See *destination address affinity persistence*.

STP (Spanning Tree Protocol)

Defined by IEEE, STP is a protocol that provides loop resolution in configurations where one or more external switches are connected in parallel with the BIG-IP system. See also *RSTP (Rapid Spanning Tree Protocol)* and *MSTP (Multiple Spanning Tree Protocol)*.

subdomain

A subdomain is a sub-section of a higher level domain. For example, **.com** is a high level domain, and **F5.com** is a subdomain within the **.com** domain.

TACACS (Terminal Access Controller Access Control System)

TACACS is an older authentication protocol common to UNIX systems. TACACS allows a remote access server to forward a user's login password to an authentication server.

TACACS+

TACACS+ is an authentication mechanism designed as a replacement for the older TACACS protocol. There is little similarity between the two protocols, however, and they are therefore not compatible.

TACACS+ authentication module

A TACACS+ authentication module is a user-created module that you implement on an LTM system to authenticate client traffic using a remote TACACS+ server.

tagged interface

A tagged interface is an interface that you assign to a VLAN in a way that causes the system to add a VLAN tag into the header of any frame passing through that interface. Tagged interfaces are used when you want to assign a single interface to multiple VLANs. See also *VLAN (virtual local area network)*.

Tcl

Tcl (Tools Command Language) is an industry-standard scripting language. On the LTM system, users use Tcl to write iRules™.

transparent node

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

trunk

A trunk is a logical group of BIG-IP system interfaces. When you create a trunk, the BIG-IP system can aggregate the corresponding links as a way to increase bandwidth.

trusted CA file

A trusted CA file is a file containing a list of certificate authorities that an authenticating system can trust when processing client requests for authentication. A trusted CA file resides on the authenticating system and is used for authenticating SSL network traffic.

Type of Service (ToS) level

The Type of Service (ToS) level is another means, in addition to the Quality of Service (QoS) level, by which network equipment can identify and treat traffic differently based on an identifier.

Universal Inspection Engine (UIE)

The Universal Inspection Engine (UIE) is a feature that offers universal persistence and universal content switching, to enhance your load balancing capabilities. The UIE contains a set of rule variables and functions for building expressions that you can specify in pool definitions and rules.

universal persistence

Universal persistence gives you the ability to persist on any string found within a packet. Also, you can directly select the pool member to which you want to persist.

user role

A user role is a type and level of access that you assign to a BIG-IP system user account. By assigning user roles, you can control the extent to which BIG-IP system administrators can view or modify the BIG-IP system configuration.

virtual address

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

virtual port

A virtual port is the port number or service name associated with one or more virtual servers managed by the LTM system. 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 an LTM system or other type of host server.

VLAN (virtual local area network)

A VLAN is a logical grouping of interfaces connected to network devices. You can use a VLAN to logically group devices that are on different network segments. Devices within a VLAN use layer 2 networking to communicate and define a broadcast domain.

VLAN group

A VLAN group is a logical container that includes two or more distinct VLANs. VLAN groups are intended for load balancing traffic in a layer 2 network, when you want to minimize the reconfiguration of hosts on that network. See also *VLAN (virtual local area network)*.

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 (virtual local area network)*.

VLAN tag

An IEEE standard, a VLAN tag is an identification number inserted into the header of a frame that indicates the VLAN to which the destination device belongs. VLAN tags are used when a single interface forwards traffic for multiple VLANs.

watchdog timer card

A watchdog timer card is a hardware device that monitors the BIG-IP system for hardware failure.

wildcard virtual server

A wildcard virtual server is a virtual server that uses an IP address of **0.0.0.0**, ***** or **"any"**. A wildcard virtual server accepts connection requests for destinations outside of the local network. Wildcard virtual servers are included only in Transparent Node Mode configurations.

WKS (well-known services)

Well-known services are protocols on ports **0** through **1023** that are widely used for certain types of data. Some examples of some well-known services (and their corresponding ports) are: HTTP (port **80**), HTTPS (port **443**), and FTP (port **20**).



Index

A

access control
 through user accounts 10-8
 to SNMP data 11-3

access levels
 and remote user accounts 10-8
 managing 10-4, 11-5

account names 10-4

account names, redundant systems 12-19

accounts
 creating and storing 10-8
 managing 10-4

actions
 at failover 12-23, 12-24
 for traps A-1

active connection statistics 11-13, 11-14

Active Directory authentication servers 1-3, 10-1, 10-8

active state
 and failback 12-4
 forcing to 12-35
 switching to 12-36

active units 12-18

active/standby mode
 and failover 12-3, 12-14
 and floating IP addresses 12-7
 and MAC addresses 12-29
 as default 12-11
 described 12-1
 returning to 12-38
 specifying 12-13

active-active mode
 and failover 12-4
 and floating IP addresses 12-7
 and MAC addresses 12-29
 described 12-1
 specifying 12-13
 updating fail-over daemon B-1

actors, trunks 8-9

address space 6-1, 7-14

admin account
 and user roles 10-3
 as default value 12-19
 configuring and managing 10-4
 described 10-2

admin account passwords 12-19

Administrator user role
 and passwords 10-6
 described 10-3, 10-4
 for account modification 10-6
 for local accounts 10-4
 for remote accounts 10-8
 for system and Web UI accounts 10-3
 for viewing accounts 10-5
 for Web UI accounts 10-3

alarm RMON group 11-12

alert system 11-7, A-1

alerts 2-3

Alternate port role 9-20

application-specific MIB files 11-1
 See also enterprise MIB files

ARP forwarding 7-18

ARP requests 12-31

ARP tables 12-15

authentication 10-1

authentication server types 10-1

authentication warnings 11-7

authentication-related traps A-4

authorization
 configuring for users 10-8
 defined 10-1
 ensuring 10-10

auto detection 12-34

auto link type 9-17

automatic failback
 and redundancy preference 12-14, 12-37
 enabling 12-37
 for active/standby mode 12-36
 for active-active mode 12-37

automatic link selection 8-7

B

backup files 12-18

Backup port role 9-20

bandwidth
 increasing 5-9, 8-1
 maximum 8-7

BCM56XXD service 12-24

BIGD service 12-24

bigdb database keys
 enabling for failback 12-37
 for redundant systems 12-36
 for SNMP 11-3

BIG-IP alert system 11-2

BIG-IP objects, SNMP 11-2

BIG-IP system information 11-3

bigpipe db command B-1

bigpipe utility 1-5, 12-21

bigtop utility 1-5

BIND forwarder server list 4-11

Blocking port state 9-21

BPDUs
 and interfaces 9-17
 defined 9-1
 maximum transmittal of 9-8
 receiving 9-18

bridge priorities 9-13

bridge protocol data units
 See BPDUs.

bridging loops
 preventing 9-1

browser versions 1-5

browser-based utility 1-5
browser-only access 4-4

C

calculations 11-13
certificates 4-5
client access
 allowing 11-2, 11-4
 configuring 11-3, 11-4
clients, SNMP 11-4
colors 12-35
command menu
 defined 3-1
 for SSH 3-3
command recommendations, lights-out 3-5
command-line configuration, lights-out 3-1
command-line utilities 1-5
commands
 for host console shell 3-4
 for out-of-band management 3-1
 See also bigpipe db command
communities
 and access levels 11-5, 11-7
 and trap destinations 11-8
community access 11-5
company-specific MIB files 11-1
ConfigSync feature 12-3
ConfigSync screen settings
 configuring 12-19
 for peer IP address 12-34
 for status 12-33
configuration data, redundant systems 12-19, 12-33
configuration files, SNMP 11-3
configuration settings, global 9-5
configuration settings, redundant systems 12-11
configuration synchronization
 and bigdb keys B-2
 and self IP addresses 12-33
 and virtual servers, SNATs 12-18
 defined 12-3
 describing 12-18
 determining need for 12-34
 for active-active mode 12-21
 problems with 12-13
configuration task summary 11-2
configuration tasks
 for active-active mode 12-11
 for LCD 2-2
 for SNMP agent 11-3
Configuration utility
 and bigdb keys B-1
 and synchronization status 12-34
 purpose of 1-5

connection mirroring
 and bigdb keys B-2
 and failback 12-36, 12-37
 described 12-28
connection processing 12-21
connection statistics 11-13, 11-15
connections
 and failover 12-16
 dropping 12-36
 failing over to peer 12-35
 managing 12-38
 preserving 12-28
 processing 12-17
 reclaiming 12-6
 relinquishing 12-5, 12-37
 See also connection mirroring
 See also failback
 See also real-time connections
console access 3-1, 10-6
console-only access 4-4
contact information 11-3
contact name 11-3
control buttons, LCD 2-1
control packets 8-6
CPU count 4-10
CPU mode 4-10
CPU use statistics 11-13, 11-17

D

data access control, SNMP 11-3
data groups 4-15
data object values, SNMP 11-1
data objects
 in MIB files 11-9
 modifying 11-5, 11-7
 See also access levels
data replication 12-3
 See also traffic replication
data synchronization
 and self IP addresses 12-33
 determining need for 12-34
 See also configuration synchronization
data validity 9-7
data, MIB files 11-11
default access levels
 assigning 11-7
 modifying 11-5
default routers 12-31
default routes
 and ARP requests 12-26
 configuring 12-16
 for management ports 4-3
default SNATs 12-22
default unit IDs 12-21

- default user roles
 - and Administrator role 10-3
 - assigning 10-9, 10-10
 - for local accounts 10-4
 - for remote accounts 10-8
- delivery failures 12-7
- Designated port role 9-20
- destination IP addresses
 - and persistence 4-16
- destinations, SNMP 11-7, 11-8
- device properties
 - viewing and configuring 4-9
- Disabled mode, for spanning tree 9-6
- Disabled port role 9-20
- display properties 2-6
- DNS (Domain Name System) 4-11
- DNS proxy services 4-11
- DNS server lookup list 4-11
- documentation, other 1-6
- dual-system configuration
 - See redundant systems
- duplex modes 5-3, 8-4
- dynamic entries 7-11

E

- edge ports 9-18
- enterprise MIB files
 - and Configuration utility 11-1
 - content of 11-10
 - defined 11-1
 - downloading 11-2, 11-9
- event notifications, SNMP 11-2
- event RMON group 11-12
- events, redundant systems 12-23
- external path cost 9-23

F

- F5-BIGIP-COMMON-MIB.txt file 11-9, A-1
- F5-BIGIP-LOCAL-MIB.txt file 11-9, 11-10
- F5-BIGIP-SYSTEM-MIB.txt file 11-9, 11-11
- failback
 - automatic 12-37
 - defined 12-5, 12-36
 - for active/standby mode 12-36
 - for active-active mode 12-37
 - initiating 12-14
- failover
 - and bigdb keys B-2
 - and floating IP addresses 12-17
 - and interfaces 12-15
 - and virtual servers 12-16, 12-17
 - defined 12-3
 - initiating 12-23, 12-26
 - triggering 12-23, 12-25
- failover actions 12-23

- failover addresses, VLANs 12-13
- failover behavior 12-3
- failover process 12-28
- failover service
 - See SOD service
- failover speed 12-29
- failover state, LCD 2-4
- failover types 12-15
- fail-safe 7-10, 12-8, 12-23
- fail-safe timeout period 12-26
- failure detections
 - for switch board components 12-23
 - for VLANs 12-26
- Fastest load balancing mode 4-14
- floating self IP address configuration
 - and VLANs 12-8
 - example of 12-7, 12-8
- floating self IP addresses
 - and default routes 12-16
 - and virtual servers 12-17
 - assigning 12-7
 - defined 6-1, 12-6
 - specifying 6-5
- floods 9-1
- flow control 5-5
- Force to Standby button 12-35
- formatting conventions 1-7
- Forward Delay option 9-7
- Forwarding port state 9-21
- frame control 5-5
- frame order 8-1
- frames 7-17
- full duplex mode 5-4

G

- gateway fail-safe 12-25
- gateway fail-safe pools 12-25
- gateway routers 12-25
- general traps A-1
- global statistics data 11-11
- graphs, SNMP 11-13
- gratuitous ARP requests 12-31
- Guest user role 10-3, 10-6

H

- half duplex mode 5-5
- hardware components 12-23
- hardware driver
 - See BCM56XXD service
- hardware-related traps A-2
- hard-wired failover 12-15
- health monitors service
 - See BIGD service
- heartbeat failures 12-23, 12-24
- Hello Time option 9-7

- help pages, LCD 2-4
- history RMON group 11-12
- Hold mode 2-2
- host console shell
 - See hostconsh shell
- host IP addresses 4-2
- host names 4-2
- hostconsh shell
 - and out-of-band protocol 3-4
 - defined 3-1
 - for menu access 3-1, 3-3
- hosts
 - and VLANs 7-5
- hosts.allow file 11-3
- HTTP request statistics 11-13, 11-16
- hub topologies 12-29

I

- icon colors 12-35
- IDs, redundant systems 12-35
- IEEE specifications 9-1
- information collection 11-2
- Information menu, LCD 2-1, 2-4
- information polling 11-2
- information validity 9-7
- information, SNMP 11-3
- instance IDs 9-13
- instance members
 - defined 9-10
 - deleting 9-15
- instances
 - allowing and creating 9-11
 - and MSTP 9-2
 - and RSTP 9-2
 - and STP 9-2
 - creating and modifying 9-12
 - defined 9-10
 - deleting 9-15
 - enabling 9-13
 - viewing and modifying 9-14
 - viewing interfaces for 9-19
- interface access methods 7-6
- interface cards 12-26
- interface costs 9-22
- interface down time 12-15
- interface media types 5-3
- interface mirroring 5-6
- interface modes 5-3
- interface naming conventions 5-2
- interface priorities 9-22
- interface properties 5-3
- interface state 5-4
- interface statistics 5-7

- interfaces
 - and frames 5-5
 - and instance settings 9-21
 - and multiple VLANs 7-6
 - and trunks 8-11
 - and VLAN traffic 5-9
 - as edge ports 9-18
 - assigning 7-5
 - blocking 9-7
 - choosing for traffic 9-22
 - configuring for spanning tree 9-16
 - defined 5-1
 - enabling and disabling 9-17
 - enabling STP on 9-10
 - for trunks 8-1, 8-4
 - monitoring 11-12
 - viewing 12-29
 - viewing per instance 9-19
- internal path cost 9-23
- internal static self IP addresses 12-6
 - See also static self IP addresses
- internet connections 12-25
- interruptions
 - detecting 12-8, 12-23
 - preventing 12-6, 12-28
- IP addresses
 - for SNMP traps 11-8
 - See also floating self IP addresses
 - See also static self IP addresses
 - specifying 11-4, 12-13
- IPv4 format 4-3
- ISO images 3-1

L

- I2 forwarding
 - See layer 2 forwarding
- LACP 8-2
- LACP control packets 8-6
- LACP packets 8-2
- LACP timeout values 8-6
- last hops 4-13
- layer 2 cache aging time 7-12
- layer 2 forwarding 7-11
- layer 2 forwarding tables 4-14
- layer 2 protocols 9-1
- LCD
 - and alerts 2-3
 - for system menu 2-4
- LCD command menu 3-3
- LCD configuration tasks 2-2
- LCD control buttons 2-1
- LCD display properties 2-6
- LCD help pages 2-4
- LCD menu mode 2-4
- LCD menus 2-1, 2-4, 4-3

-
- LCD panel 2-2
 - LCD screen alerts 2-3
 - LCD screens 2-2
 - LDAP authentication servers 1-3, 10-1, 10-8
 - LDAP database 10-4
 - learning port state 9-21
 - LED activity 2-4
 - legacy bridges
 - and BPDUs 9-18
 - and MSTP 9-2
 - and STP 9-2
 - license-related traps A-2
 - lights-out management 3-1
 - lights-out system 3-1
 - link aggregation
 - and path costs 9-23
 - defined 5-9, 8-1
 - Link Aggregation Control Protocol
 - See LACP
 - link aggregation policy 8-6
 - link selection policy 8-6
 - link types, for STP 9-17
 - load balancing 1-5
 - local application traffic 11-10
 - local traffic management information 11-9
 - local user accounts
 - creating 10-5
 - managing 10-1, 10-4
 - local user-account database 10-8
 - locked-up system 3-1
 - logging-related traps A-5
 - long timeout value 8-6
 - loopback interface 11-4
 - loops 9-1
 - lossy networks 12-29
- M**
- MAC address changes 12-31
 - MAC addresses
 - assigning 6-2
 - finding 7-9
 - for trunks 8-1
 - for VLANs 7-6, 12-29
 - sharing 12-2
 - viewing 12-30
 - MAC masquerade addresses
 - defined 7-8
 - sharing 12-2, 12-29
 - specifying 7-17
 - support of 12-13
 - machine location 11-4
 - maintenance mode 4-13
 - management information base
 - See also MIB
 - See MIB-II MIB
 - management interface 5-2
 - management port 3-3
 - management port route 4-2
 - management options, lights-out 3-5
 - manual failback 12-36
 - Maximum Age option 9-7
 - MCPD service 12-24
 - media access control addresses
 - See MAC addresses
 - media properties
 - monitoring 8-6
 - viewing 5-3
 - media speed 8-4
 - media type 5-4
 - media types 5-3
 - members, MSTP 9-15
 - memory use 4-14
 - memory use statistics 11-13, 11-14
 - Menu mode, LCD 2-2
 - menus
 - See LCD menus
 - messages
 - and VLANs 7-6
 - See also notification messages
 - See traps
 - messaging service
 - See MCPD service
 - metrics collection 11-13
 - MGMT interface 5-2
 - MGMT port 3-3
 - MIB 11-9
 - and device management 11-1
 - defined 11-1
 - See also MIB-II MIB
 - MIB file contents 11-10
 - MIB file locations 11-1
 - MIB file types 11-9
 - MIB files
 - defined 11-9
 - described 11-1
 - downloading 11-2
 - MIB files location 11-9
 - MIB-II MIB 11-1
 - MIB-II objects 11-11
 - mirroring 5-6
 - See connection mirroring
 - modes, duplex 5-3, 8-4
 - monitoring 1-5, 12-3
 - MSTP
 - and VLANs 9-3, 9-11
 - defined 9-2
 - MSTP Configuration Name setting 9-8
 - MSTP Configuration Revision setting 9-8
 - MSTP Maximum Hop Number setting 9-9
 - MSTP mode, for spanning tree 9-6
 - MSTP regions 9-3, 9-8

MTUs 4-13, 4-14
multilayer switches 7-1
Multiple Spanning Tree Protocol
 See MSTP.

N

name resolution 4-11
named service 4-11
naming conventions, interfaces 5-2
netmasks 6-4
Net-SNMP 11-1
network boot 3-1, 4-10
network connections 12-3, 12-25
network failover 12-15
network failure 9-1
network failure detection 12-25
network information 11-11
network monitoring 12-8, 12-23
Network Time Protocol
 See NTP
network traffic
 monitoring 12-26
 reclaiming 12-36
network-related traps A-4
new connection statistics 11-13, 11-15
No Access user role 10-3
noarp mode 12-8
notification events 11-7
notification messages 11-2, 11-7
 See also traps
notification types A-1
notifications 11-10
notifications, SNMP 11-10
NOTIFICATION-TYPE designation 11-10
NTP 4-10

O

object data, SNMP 11-2
object ID definitions, RMON 11-12
object presentation 11-1
object values, SNMP 11-1
OIDs 11-13
online help types 1-6
opaque mode 7-16
Operator user role 10-3, 10-6
optional tasks, redundant systems 12-10
Options menu 2-1
out-of-band management commands 3-1

P

p2p link type 9-17
packet collisions 12-30
packet floods 9-1
packet forwarding 4-15

packets
 and access to VLANs 7-6
Pass Through mode, for spanning tree 9-6
passwords
 assigning 10-4
 changing 4-4, 10-6
 defined 10-1
 for synchronization 12-19
path costs 9-22
paths
 blocking and clearing 9-1
Pause frames 5-5
peer communication 12-6
peer unit 12-20
peer unit monitoring 12-3
peer, trunks 8-9
performance improvement, redundant systems 12-4
performance improvements 9-2
performance metrics, SNMP 11-2, 11-13
permissions 10-4
persistence information 12-28
persistence table 4-15
platform information 11-11
platform properties 4-2
pool member thresholds 12-25
pool traffic 12-25
port lockdown 6-5
port numbers 11-8
port roles 9-20
port states 9-21
ports
 and VLANs 7-6
ports, lights-out 3-1
preferred redundancy states 12-14
primary failover addresses 12-12
primary failover pair 12-6
properties, local traffic 4-13
properties, redundant systems 12-35
protocol degradation legacy bridges
 and protocol degradation 9-3
protocol versions 9-3
protocols
 accepting traffic from 6-5
proxied ARP requests 7-18
proxy IP addresses 4-15
PXE installation 4-3

Q

quiet boot 4-10

R

RADIUS authentication servers 1-3, 10-1, 10-8
Rapid Spanning Tree Protocol
 See RSTP.
rate statistics 11-16

- read/write access level 11-5, 11-7
 - read-only access level 11-5, 11-7
 - real-time connections 12-28
 - reaper high-water mark 4-14
 - reaper low-water mark 4-14
 - rebooting
 - amd memory use 4-10
 - redundancy modes 12-1, 12-13
 - redundancy state 12-32, 12-36
 - redundancy state preferences
 - setting 12-37
 - viewing 12-32
 - redundancy state preferences, listed 12-14
 - redundancy states
 - switching 12-35
 - viewing 12-32
 - redundant paths
 - blocking and clearing 9-1
 - redundant system customization 12-36
 - redundant system features 12-1
 - redundant system modes 12-38
 - redundant systems
 - advanced features of 12-38
 - and active-active mode 12-21
 - and server routing 12-16
 - defined 12-1
 - maintaining 12-32
 - See also connection mirroring
 - See also fail-safe
 - See configuration synchronization
 - setting up 12-11
 - sharing MAC addresses 7-9, 12-30
 - redundant-pair designation 12-11
 - reference links 8-2
 - regions 9-3
 - remote account databases 10-8
 - remote authentication server types 10-1
 - remote authentication servers 10-8
 - remote lights-out management 3-3
 - Remote Network Monitoring
 - See RMON implementation
 - remote server types 10-8
 - remote SSH access, setting up 3-4
 - remote system management 3-1, 4-3, 11-4
 - remote user accounts
 - assigning user roles to 10-9
 - changing 10-10
 - creating and storing 10-8
 - remote user roles
 - assigning 10-9
 - changing 10-10
 - replication 12-3
 - See traffic replication
 - reset option, lights-out 3-6
 - responses
 - processing 12-16
 - receiving 12-7
 - RFC 826 12-31
 - RMON groups 11-12
 - RMON implementation 11-12
 - RMON-MIB.txt file 11-12
 - root account
 - and user roles 10-3
 - modifying 10-6
 - root account requirements 10-2
 - root bridges 9-13
 - Root port role 9-20
 - Rotate mode, LCD 2-2
 - routers
 - and noarp mode 12-8
 - notifying 12-31
 - routes
 - blocking 5-9
 - routing tables 12-16
 - RSTP
 - and legacy bridges 9-2
 - defined 9-2
 - RSTP mode, for spanning tree 9-6
- ## S
- SCCP utility 3-1, 3-3
 - screen alerts 2-3
 - Screen menu, LCD 2-1
 - secondary failover addresses 12-13
 - secondary failover pair 12-6
 - secure hub topologies 12-29
 - Secure Sockets Layer connections 12-28
 - security, adding 10-1
 - self IP addresses 6-3
 - self IP addresses
 - and unit IDs 12-21, 12-22
 - assigning 7-15, 7-19
 - assigning to VLANs 7-4
 - defined 6-1, 12-6
 - for SNATs 6-2
 - managing 6-6
 - See also floating self IP addresses
 - See also static self IP addresses
 - specifying 12-13
 - serial console 3-3
 - server connections 12-17
 - server responses 12-16
 - servers 12-16
 - service interruptions 12-1
 - service, returning to 12-14
 - services
 - monitoring 12-23, 12-24
 - See also system services

- Setup utility
 - and user accounts 10-1
 - as prerequisite task 12-8
- shared floating IP addresses 12-7
 - See also floating self IP addresses
- shared link type 9-17
- shared MAC addresses 4-15
- short timeout value 8-6
- shutdowns 3-6
- simultaneous connection processing 12-21
- simultaneous traffic management 12-4
- SNAT association 12-5
- SNAT automapping 6-2
- SNAT information 11-9
- SNAT packet forwarding 4-15
- SNAT pools 6-2
- SNAT traffic 12-7
- SNATs
 - and active-active mode 12-21
 - and configuration data 12-18
 - and connection mirroring 12-28
 - and route configuration 12-16
 - and unit IDs 12-21, 12-22
 - defined 6-2
 - managing connections for 12-38
 - use of 7-13
- SNMP
 - and syslog 11-9
 - configuring 11-7
 - in the Configuration utility 11-4
 - See also SNMP managers
 - See SNMP agent
- SNMP agent
 - access to 11-4
 - and trap notifications A-1
 - configuring 11-2
 - defined 11-1
- SNMP client 11-4
- SNMP commands
 - for collecting statistics 11-13
 - using 11-2, 11-10
- SNMP configuration files 11-3
- SNMP data access control 11-3
- SNMP manager functions 11-2
- SNMP managers
 - as trap destinations 11-8
 - defined 11-1
- SNMP MIB
 - See MIB
- SNMP MIB files
 - See MIB files
- SNMP object data 11-2
- SNMP tasks 11-1, 11-2
- SNMP traps
 - handling 11-2
 - See also traps
- SNMP user access 11-6
- SNMP user names 11-3
- SNMP users 11-5
- snmpd service 11-3
- snmpd.conf files
 - and access levels 11-6
 - and SNMP agent 11-3
 - for trap configuration 11-8
- snmpget command 11-13
- SOD service 12-24
- software upgrades 4-3
- source IP addresses 12-6
- source network address translations
 - See SNATs
- spanning tree instances
 - creating and modifying 9-12
 - defined 9-10
 - deleting 9-15
 - enabling 9-13
 - viewing and modifying 9-14
- Spanning Tree Protocol
 - See STP
- spanning tree protocols
 - and trunks 8-5
 - changing 9-7
 - resetting 9-18
- spanning tree regions
 - See MSTP regions
- spanning trees 9-2
- SSH 4-3, 4-5
- SSH access 10-6
- SSH clients 1-5, 3-3
- SSL certificates 4-5
- SSL connections 12-28
- standby state
 - changing 12-4
 - forcing to 12-35
- state duplication 12-28
- state preference 12-37
- state transitioning 12-31
- states 12-35
 - See also active state
 - See also standby state
- states, redundant systems 12-36
- static entries 7-11
- static self IP addresses
 - assigning 12-6
 - defined 6-1, 12-6
 - for data synchronization 12-33
 - specifying 12-13
- statistical data 11-11
- statistics 5-7
- statistics RMON group 11-12
- statistics, SNMP 11-11

- status
 - and synchronization 12-34
 - communicating 12-6
 - viewing 12-35
 - status messages 12-34
 - status monitoring 12-3
 - STP 9-2
 - STP mode, for spanning tree 9-6
 - style conventions 1-7
 - support account
 - enabling 4-4
 - requiring 10-2
 - switch board monitoring 12-23
 - switch card control process
 - See SCCP utility
 - See BCM56XXD service
 - switch interoperability 12-29
 - SYN Check 4-14
 - synchronization
 - determining need for 12-34
 - See also configuration synchronization
 - synchronization buttons 12-19
 - synchronization frequency 12-19
 - synchronization methods 12-18
 - synchronization permissions 12-19
 - synchronization status
 - and active units 12-33
 - and auto detection 12-34
 - and Configuration utility 12-19, 12-32
 - and status icon 12-35
 - enabling global view 12-18
 - synchronization status display 12-20
 - synchronization status messages 12-34
 - system access control 10-2, 10-8
 - system account management 10-4
 - system data, SNMP 11-10
 - system fail-safe
 - and bigdb keys B-2
 - and hardware components 12-23
 - system information
 - configuring 11-3
 - polling for 11-10
 - system information and statistics 2-5
 - system interface monitoring 11-12
 - system location 11-3
 - System maintenance account types 10-2
 - system maintenance accounts 10-1
 - System menu 2-1
 - system monitoring 12-8, 12-23
 - system objects, SNMP 11-2
 - system operation, SNMP A-1
 - system service failures 12-24
 - system services
 - monitoring 12-23, 12-24
 - restarting 12-26
 - system status, redundant systems 12-6
 - system time 4-10
 - system unavailability 12-36
 - system upgrades 4-3
- T**
- tagged interfaces, defined 7-6
 - tags
 - embedding in packet headers 7-6
 - See VLAN tags
 - task summary
 - for redundant system configuration 12-9
 - for SNMP 11-2
 - tasks, for active-active mode 12-21
 - TCP/IP connections 12-3
 - thresholds 12-25
 - throughput rate statistics 11-13, 11-16
 - time 4-10
 - time interval, changing 9-7
 - time zones 4-2
 - timeout periods, fail-safe 12-26
 - timer values 9-7
 - TMM service 12-24
 - TMOS-related traps A-3
 - traffic
 - blocking 9-7
 - receiving 12-7
 - reclaiming 12-6, 12-36
 - relinquishing 12-5
 - restricting through tagged interfaces 7-6
 - restricting through untagged interfaces 7-6
 - traffic distribution 8-1
 - traffic interruptions 12-6, 12-7
 - traffic management operating system 3-1
 - traffic management service
 - See TMM service
 - traffic processing 12-18
 - traffic replication 12-3
 - translucent mode 7-16
 - Transmit Hold Count option 9-8
 - transparency mode 7-14, 7-16
 - transparent mode 7-16
 - trap destinations
 - configuring 11-3
 - setting 11-7, 11-8
 - trap information storage 11-3
 - trap locations 11-9, 11-10
 - trap names A-1
 - traps
 - and recommended actions A-1
 - configuring 11-3
 - defined 11-2, 11-7
 - handling 11-2
 - identifying 11-10
 - See also notification types

- tree structure 11-1
- trunk lists 8-9
- trunk properties 8-9
- trunks 5-9
 - and spanning tree protocols 8-5
 - and VLANs 7-5, 8-10
 - creating 8-3
 - defined 8-1
 - deleting 8-10
 - enabling and disabling 9-17
- trunks names 8-4

U

- UCD-SNMP 11-1
- unit IDs
 - assigning 12-36
 - associating 12-21
 - configuring 4-2
 - determining 12-35
 - purpose of 12-3
 - viewing 12-21
- units
 - halting 2-2
 - powering down 2-2, 2-3
 - powering on 2-2
 - rebooting 2-3
 - resetting 3-1
- untagged interfaces 7-6, 8-4
- untagged interfaces, defined 7-6
- upgrades 4-3
- user account configuration tools 10-2
- user account location 10-1
- user account maintenance 1-3
- user account types 10-1
- user accounts
 - and tools 10-2
 - changing 10-10
 - creating 10-1, 10-8
 - managing 10-4
 - purpose of 10-1
 - See also remote user accounts
 - See local user accounts
 - storing 10-8
- user administration 4-2
- user authorization 10-8
- user name storage, SNMP 11-3
- user names, general 10-4
- user names, redundant systems 12-19
- user role categories 10-2
- user role requirements 10-9
- user roles
 - and account creation 10-4
 - and Web UI accounts 10-1
 - as authorization 10-2
 - assigning 10-8

- changing 10-6, 10-10
- defined 10-1, 10-2
- described 10-3
- for remote accounts 10-8

users

- See also user accounts
- See SNMP users

V

- virtual server association 12-5
- virtual server connections 12-16
- virtual server information 11-9
- virtual server traffic 12-7
- virtual servers
 - and active-active configuration 12-5
 - and active-active mode 12-21
 - and configuration data 12-18
 - and connection mirroring 12-28
 - and destination addresses 4-14
 - and failover 12-17
 - and unit IDs 12-21, 12-22
 - managing connections for 12-38
- VLAN access methods 7-6
- VLAN configuration 7-3
- VLAN fail-safe 12-26
- VLAN group IDs
 - assigning 7-14
 - defined 7-16
- VLAN group names 7-15
- VLAN groups
 - and self IP addresses 6-5
 - creating 7-15
 - defined 7-1, 7-13
- VLAN IDs 7-6
- VLAN members
 - adding 9-10
 - removing 9-15
- VLAN names 7-14, 12-26
- VLAN tags 7-6, 9-14
- VLAN traffic 5-9
- VLAN-keyed connections 4-13
- VLANs
 - adding as instance members 9-12, 9-14
 - and connections 4-13
 - and failover addresses 12-6
 - and fail-safe 12-26
 - and failure detection 12-26
 - and interfaces 5-1
 - and MAC addresses 12-29
 - and MSTP 9-2
 - and self IP addresses 6-4
 - and STP 9-2
 - and trunks 8-5, 8-10
 - benefits of 7-1
 - creating 7-4

defined 7-1
example of 7-3

W

warnings 11-7
web certificates 4-5
Web UI account types 10-2
Web UI accounts
 defined 10-1
 managing 10-4
web-based utility
 See browser-based utility

X

X button 2-4