BIG-IP® Link Controller Solutions Guide

version 4.5.10

MAN-0096-00

Product Version

This manual applies to version 4.5.10 of the BIG-IP® Link Controller.

Legal Notices

Copyright

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described 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, GLOBAL-SITE, SEE-IT, EDGE-FX, FireGuard, Internet Control Architecture, IP Application Switch, Control Your World, PACKET VELOCITY, SYN Check, uRoam, and FirePass 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.

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 contains software based on oprofile, which is protected under the GNU Public License.

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.

This product includes software developed by Eric Young.

This product includes Malloc library software developed by Mark Moraes. (© 1988, 1989, 1993, University of Toronto).

This product includes open SSL software developed by Eric Young (eay@cryptsoft.com), (© 1995-1998).

Table of Contents

Introduction

1 Introducing the Link Controller

2 Configuring Links for Simple ISP Load Balancing

3 Configuring Cost-Based ISP Load Balancing

4 Working with Link Controllers in a 3-DNS Sync Group

Introduction

- Getting started
- Using the Administrator Kit
- Finding additional help and technical support resources
- Learning more about the BIG-IP product family

Getting started

Before you start configuring the BIG-IP Link Controller, we recommend that you browse this guide to find the solution that most closely addresses your needs. Briefly review the basic configuration tasks and the few pieces of information, such as IP addresses and host names, that you should gather in preparation for completing the tasks.

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

Choosing a configuration tool

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

The Setup utility

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

The Configuration utility

The Configuration utility is a web-based application that you use to configure and monitor the setup on the Link Controller. Once you complete the installation instructions described in this guide, you can use the Configuration utility to perform the configuration steps necessary for your chosen load balancing solution. In the Configuration utility, you can also monitor current system performance, and download administrative tools such as the SNMP MIB or the SSH client. The Configuration utility is best viewed with the following browsers: Netscape® Navigator version 4.7x, or Microsoft® Internet Explorer version 5.0, 5.5, or 6.0.

The bigpipe and bigtop command line utilities

The **bigpipe**™ utility is the command line counterpart to the Configuration utility. Using **bigpipe** commands, you can configure virtual servers, open ports to network traffic, and configure a wide variety of features. To monitor the Link Controller, you can use certain **bigpipe** commands, or you can use the **bigtop**™ utility, which provides real-time system monitoring. You can use the command line utilities directly on the Link Controller console, or

you can run commands using a remote shell, such as the SSH client (encrypted communications only), or a Telnet client (if you are restricted by cryptography export laws). For detailed information about the command line syntax, see Appendix A in the ***BIG-IP Reference Guide***.

Using the Administrator Kit

The Link Controller Administrator Kit provides all of the documentation you need to work with the BIG-IP Link Controller. The Link Controller includes the following printed documentation:

- **Configuration Worksheet**
  This worksheet provides you with a place to plan the basic configuration for the BIG-IP system.

The following guides are available in PDF format from the CD-ROM provided with the BIG-IP system. These guides are also available from the home screen when you first log in to the Configuration utility on the BIG-IP system.

- **Platform Guide**
  This guide includes information about the BIG-IP unit. It also contains important environmental warnings.
- **BIG-IP Link Controller Solutions Guide**
  This guide provides examples of common link load balancing solutions using the Link Controller. Before you begin configuring the Link Controller, we recommend that you browse this guide to find the load balancing solution that works best for you.
- **BIG-IP Reference Guide**
  This guide provides detailed configuration information for all BIG-IP systems, including the Link Controller. It also provides syntax information for **bigpipe** commands, other command line utilities, configuration files, system utilities, and monitoring and administration information.

Stylistic conventions

To help you easily identify and understand important information, this section describes the stylistic conventions used in our documentation.

Using the solution examples

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

Identifying new terms

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

Identifying references to objects, names, and commands

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

Identifying references to other documents

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

Identifying command syntax

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

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

or

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

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

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

Table Intro.1 *Command line syntax conventions*

Item in text	Description
[ ]	Indicates that syntax inside the brackets is optional.
...	Indicates that you can type a series of items.

Table Intro.1 *Command line syntax conventions*

Finding additional help and technical support resources

You can find additional technical information about this product using the following resources:

- **Release notes**
  Release notes for the current version of this product are available from the product web server home page, and are also available on the technical support site. The release notes contain the latest information for the current version, including a list of new features and enhancements, a list of fixes, and, in some cases, a list of known issues.
- **Online help**
  You can find help online in three different locations:
  - The web server on the product has PDF versions of the guides included in the Administrator Kit.
  - The web-based Configuration utility has online help for each screen. Simply click the **Help** button.
  - Individual **bigpipe** commands have online help, including command syntax and examples, in standard UNIX™ man page format. Simply type the command followed by the word **help**, and the Link Controller displays the syntax and usage associated with the command.
- **Third-party documentation for software add-ons**
  The Welcome screen in the Configuration utility and the resource CD-ROM both contain online documentation for all third-party software, such as Namesurfer.
- **Technical support through the World Wide Web**
  The F5 Networks Technical Support web site, **http://tech.f5.com**, provides the latest technical notes, answers to frequently asked questions, updates for Administrator Kit guides (in PDF format) and Release Notes, and the Ask F5 natural language question and answer engine. You need a user ID and password to access this site. If you do not already have a user ID and password, you can register at **http://tech.f5.com** and get immediate access to the site.

◆ Note

All references to hardware platforms in this guide refer specifically to systems supplied by F5 Networks, Inc. If your hardware was supplied by another vendor and you have hardware-related questions, please refer to the documentation from that vendor.

Learning more about the BIG-IP product family

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

- **The BIG-IP system**
  The BIG-IP HA, HA+, and 5000 software provides the full suite of local area load balancing functionality. The BIG-IP unit also has an optional 3-DNS software module which supports wide-area load balancing.
- **The BIG-IP Link Controller**
  The complete version of the BIG-IP software provides the full suite of local area load balancing functionality. The BIG-IP unit also has an optional 3-DNS software module which supports wide-area load balancing.
- **The BIG-IP e-Commerce Controller**
  The BIG-IP e-Commerce Controller uses SSL acceleration technology to increase the speed and reliability of the secure connections that drive e-commerce sites.
- **The BIG-IP special purpose products**
  The special purpose BIG-IP system provides the ability to choose from three different BIG-IP feature sets. When you run the Setup utility, you specify one of three types:
  - **The BIG-IP Load Balancer**
    The BIG-IP Load Balancer provides basic load balancing features.
  - **The BIG-IP FireGuard**
    The BIG-IP FireGuard provides load balancing features that maximize the efficiency and performance of a group of firewalls.
  - **The BIG-IP Cache Controller**
    The BIG-IP Cache Controller uses content-aware traffic direction to maximize the efficiency and performance of a group of cache servers.

1

Introducing the Link Controller

- What is the BIG-IP Link Controller?

What is the BIG-IP Link Controller?

The BIG-IP Link Controller is a dedicated IP Application Switch that manages bi-directional traffic to a site across multiple links, regardless of connection type or provider. The Link Controller provides granular traffic control for Internet gateways, so that users can define how traffic is distributed across their links in a way that meets their business priorities. The Link Controller transparently monitors the availability and health of links to optimally direct traffic across the best available link.

The Link Controller includes the following features:

- Dynamic load balancing, based on the following link attributes:
  - Total available bandwidth of the link
  - The costs of purchased incremental bandwidth segments
  - Inbound link capacity and resource limits
  - Outbound link capacity and resource limits
- Router monitoring, to ensure high availability and continuous uptime
- Internet Link Evaluator, to view current Internet connectivity metrics

◆ Note

The Link Controller is also available as a module on the BIG-IP system.

Understanding link load balancing

Link load balancing is defined as managing traffic across multiple Internet or WAN gateways. Link load balancing ensures high availability in the network, and improves the performance of a web site or data center. Link load balancing provides a method by which an organization can reliably manage a multi-homed network. A ***multi-homed network*** is composed of one or more data centers that have more than one gateway, or link, to the Internet.

As enterprises increase their reliance on the Internet for delivering mission-critical applications and services, using only one link and ISP provider to access the public network represents a single point of failure. Link Load balancing, with the BIG-IP Link Controller, removes the risk of this single point of failure by enabling enterprises to control and monitor multiple links for their Internet connectivity.

The BIG-IP Link Controller:

- Guarantees reliable network connections and eliminates downtime by detecting any type of connection outage, and transparently directing traffic away from malfunctioning or unavailable links.
- Distributes traffic to optimize the capacity of each connection by monitoring line throughput so that links do not become over-saturated.

- Increases user and site performance. The BIG-IP Link Controller measures and directs users over the best performing link to increase end user and site response times.
- Directs traffic over the least expensive link. Administrators can define the price of links and tiered pricing schemes. The BIG-IP Link Controller can direct traffic to the least expensive connection, lowering bandwidth costs.
- Traffic control to match business priorities. Organizations can define traffic policies to direct traffic over specified connections.
- Provides the Internet Link Evaluator, so that you can evaluate link performance, over time, to help you choose the Internet service providers that can best serve your customers.
- Provides integrated firewall load balancing, and enables customers to double their firewall performance through intelligent traffic management to redundant firewall configurations.

Link management

With the Link Controller, you can manage both inbound and outbound traffic over multiple links. You can distribute traffic based on performance, bandwidth cost, and bandwidth availability. The metrics you can specify are the limits on bandwidth usage, and the pricing structure of your purchased bandwidth. When you specify the limits and pricing metrics for your links, the Link Controller then load balances the links based on those metrics.

Link monitoring

You can monitor several aspects of your managed links using the following tools:

- **Link Statistics**
  The Link Statistics screen, in the Configuration utility, displays information about the status, bandwidth usage, bandwidth limits, and bandwidth costs for each of the links managed by the Link Controller.
  For more information on link statistics, review Chapter 16, *Working with Link Configuration*, in the ***BIG-IP Reference Guide***.
- **Internet Link Evaluator**
  The Internet Link Evaluator shows the average round trip times, completion rates, and router hops over the links managed by the Link Controller for any of the seven continents. You can use the Internet Link Evaluator to determine which links best serve a particular global region. You can also use the Internet Link Evaluator to compare link performance to a particular global region.
  For more information on the Link Evaluator, review Chapter 15, *Internet Link Evaluator*, in the ***BIG-IP Reference Guide***.

- **Transparent monitoring**
  Transparent monitoring provides the health status of the routers for the managed links.

2

Configuring Links for Simple ISP Load Balancing

- Introducing simple ISP load balancing
- Configuring ISP load balancing
- Additional configuration options

Introducing simple ISP load balancing

You can configure the Link Controller to provide high availability for incoming and outgoing traffic with multiple Internet service providers (ISPs). This configuration eliminates the possibility that if one ISP connection fails, your entire web site or Internet connectivity fails.

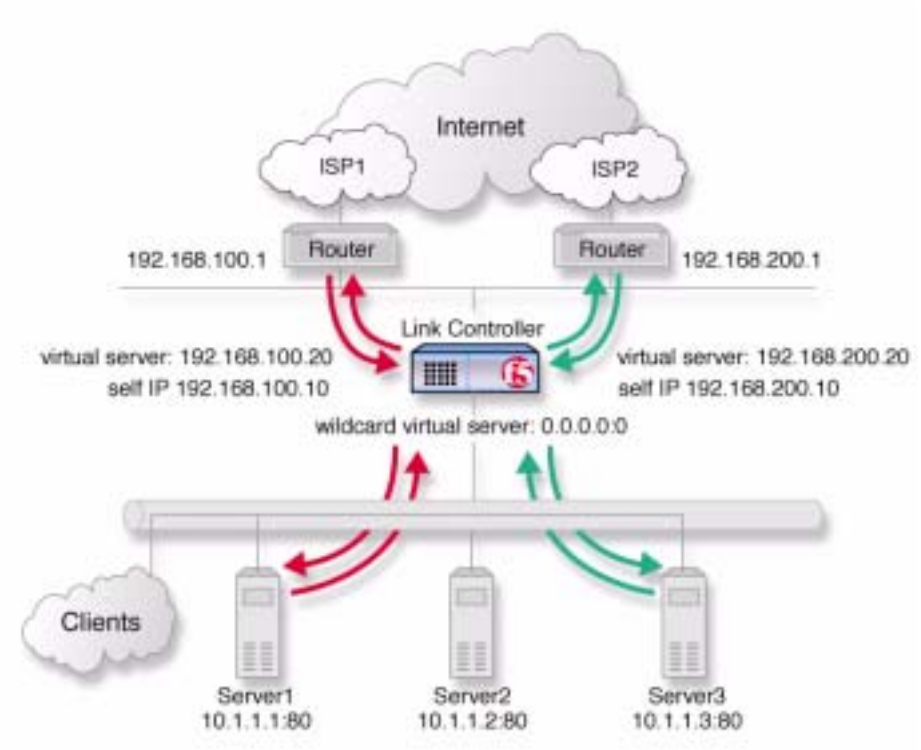

Figure 2.1 *An example of simple ISP load balancing.*

◆ Note

This type of configuration assumes that you have completed the base configuration created by the Setup utility. For more information, see the ***BIG-IP Reference Guide****, Chapter 2,* ***Using the Setup Utility****. To use this configuration, you must configure at least three VLANs when you create the initial configuration: one VLAN for each ISP, and one VLAN for the internal network.*
The IP addresses used in this example are for demonstration only. You should substitute IP addresses appropriate for your network.

Configuring ISP load balancing

When you set up ISP, or link, load balancing, you have several tasks to complete on the Link Controller:

- **Configure the links**
  Complete the following tasks to configure the links.
  - Verify that the default gateway pool that contains the IP address of each ISP, or link, is configured correctly.
  - Add the links to the configuration.
  - Create transparent monitors to verify that the path to or through ISP is available.
- **Create two load balancing pools**
  You must define one pool that load balances the content servers. The other pool, the **default_gateway_pool**, is created when you add the IP addresses of the links while setting the default gateways in the Setup utility.
- **Configure virtual servers**
  You need to configure a virtual server in the network of the link for each ISP to load balance inbound connections across the servers. You also configure one wildcard virtual server (**0.0.0.0:0**) to load balance outbound connections across the routers.
- **Add a wide IP for inbound load balancing**
  Add a wide IP to handle inbound DNS requests for each pair of virtual servers you add for each link.
- **Manage links**
  Additional configuration options are available for each link.

Defining the pools for an additional Internet connection

First, define one pool that load balances the content servers, and one pool to load balance the routers. Figure 2.2, on page 2-3 is an example of how the network devices and servers are grouped into pools.

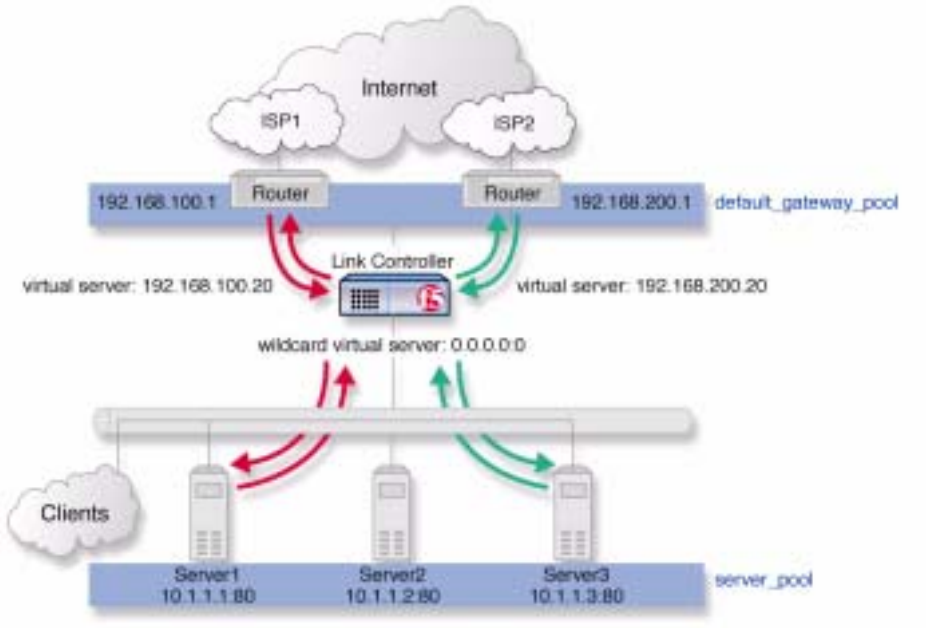

Figure 2.2 *The pools required for link load balancing*

To create the inbound load balancing pool using the Configuration utility

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

   Configuration notes

   For the example in Figure 2.1:

   Create the pool ***server_pool*** *containing the members* ***10.1.1.1:80*** ***10.1.1.2:80****, and* ***10.1.1.3:80****.*

To create the default gateway pool using the Configuration utility

If you configured more than one default gateway in the Setup utility, the Link Controller already created a **default_gateway_pool** pool. You can skip this step. If you do not have a **default_gateway_pool**, create one by completing the following task.

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

   Configuration notes

   For the example in Figure 2.1:

   *Create the pool **default_gateway_pool** containing the router inside addresses **192.168.100.1:0** and **192.168.200.1:0**.*

Defining the virtual servers for an additional Internet connection

After you create the pools, you configure the virtual servers, one for each link that load balances inbound connections across the servers. You also configure one wildcard virtual server to load balance outbound connections across the routers. Each of the virtual servers you create references either the **default_gateway_pool** or the **server_pool** in the configuration.

To define the virtual servers for inbound traffic using the Configuration utility

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

   Configuration notes

   For the example in Figure 2.1:

   *Note that you must create a virtual server for each link. For this example, create the virtual servers **192.168.100.20:80** and **192.168.200.20:80**, and use pool **server_pool**.*

To define a wildcard virtual server for outbound traffic using the Configuration utility

1. In the navigation pane, click **Virtual Servers**.
   The Virtual Servers Screen opens.
2. Click the **Add** button.
   The Add Virtual Server screen opens.
3. For a wildcard virtual server, use the following settings:
   - For the virtual server address, type **0.0.0.0**.
   - For the service, type **0** or select **any** from the list.
   - For the VLAN, select **All**.
   - For the pool, select **default_gateway_pool**.
4. Click **Done**.

   Click **Help** for details on the settings on the Add Virtual Server screens.

Setting the default gateway pool

If a Link Controller does not have a predefined route for network traffic, the unit automatically sends traffic to the pool that you define as the default gateway pool. Think of the default gateway pool as a pool of default routes.

In a Link Controller configuration, the default gateway pool must contain two or more gateway IP addresses, or ISPs. If a gateway in the default gateway pool becomes inactive, existing connections through the inactive gateway are routed through another gateway in the default gateway pool.

◆ Note

*If you configure more than one default gateway in the Setup utility, the Link Controller automatically creates a **default_gateway_pool** pool.*

To set the default gateway pool from the Configuration utility

1. In the navigation pane, click **System**.
   The System screen opens.
2. Click the Properties tab.
   The Properties screen opens.
3. From the **Default Gateway Pool** list, select the pool that contains the internal IP addresses of the gateway routers.
   In the example in this document, this is the **default_gateway_pool** pool.
4. Click **Apply**.

WARNING

Default gateway IP addresses must have a corresponding self IP address/netmask combination on the Link Controller.

Using SNAT automap for outbound traffic

Secure network address translation (SNAT) automap is automatically configured for outbound traffic so that clients receive replies through the same ISP that their requests originated from. Figure 2.3 is an example of the SNAT automap configuration for link load balancing.

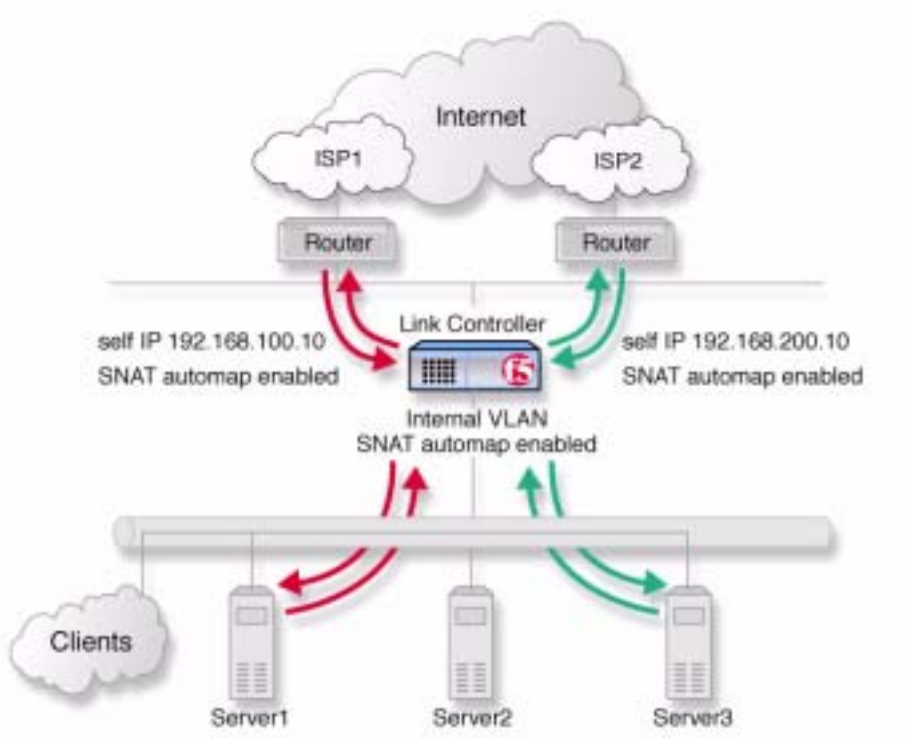

Figure 2.3 *SNAT automap feature on self IP addresses and on internal VLANs.*

Adding a wide IP for inbound load balancing

To complete the link load balancing configuration, you must configure a wide IP and wide IP pool for each pair of virtual servers you created for each link. Each wide IP in your configuration has a pool of virtual servers that the Link Controller load balances incoming DNS requests to. The wide

IP pool is made up of only virtual servers managed by the Link Controller. When you configure the wide IP pool, you specify the load balancing methods that the Link Controller applies to the incoming DNS requests.

To add a new wide IP for inbound load balancing using the Configuration utility

1. In the navigation pane, click **Link Configuration**, and then click **Inbound LB**.
   The Wide IPs list screen opens.
2. Click the **Add** button.
   The Define Wide IP (Step 1 of 2) screen opens.
3. Add the wide IP name and port settings, and click **Next**.
   The Define Wide IP (Step 2 of 2) screen opens.
4. In the **Available** list, click the virtual servers that you want to add to the pool, and click the Add (-->>) button.
   The virtual servers become part of the **Members** list for the wide IP.
5. Click **Finish**.
   The wide IP is added to your configuration.

Configuring transparent monitors for the links

When you create the default gateway pool, the Link Controller automatically creates simple ICMP monitors that check to make sure the IP addresses in the default gateway pool are available to the Link Controller. In addition to the default ICMP monitors, you can configure transparent monitors that verify the path taken by traffic through each link. You can use transparent monitors to check the availability of a device in an ISP network or on the Internet.

To configure a transparent monitor using the Configuration utility

1. In the navigation pane, expand the **Link Configuration** item, and then click **Links**.
   The Link List screen opens.
2. Click the name of the link that you want to modify.
   The Link Properties tab opens.
3. Click the Link Monitor tab.
   The Monitor screen opens.
4. Type in the name of your monitor, for example **LinkMonitor** (it must be different from the monitor template name), and select the **tcp_echo** monitor template.
5. Click the **Next** button.
   The Configure Basic Properties screen opens. The default value for **interval** and **timeout** should be sufficient.

6. Check the **Transparent** box, and click the **Next** button.
   The Configure Destination IP and Service (Alias) screen opens.
7. In the **Destination IP** box, type the IP address of a network device on the other side of the router that you want to monitor.
   - If you are monitoring a device at the ISP, you should contact the ISP for more information about how to configure this monitor.
   - If you are monitoring a device in an ISP network, you must configure a separate monitor, with a specific IP and port combination for the device, for each ISP network.
   - If you are monitoring a device on the public Internet, you can create one monitor that is applicable for all links.
8. In the **Destination Service** box, type the service number you want to monitor. For example, if Telnet is enabled on the destination device, type **23**.
9. When you have finished configuring the monitor, click **Apply**.

◆ Note

You can also set up ICMP transparent monitors. When your ISP does not allow TCP monitoring, use ICMP monitoring instead.

To associate the monitor with the members of the default gateway pool

After you create the monitors to check the availability of the links, you must associate the monitors with the routers in the default gateway pool.

1. In the navigation pane, click **Monitors**.
   The Monitors screen opens.
2. Click the Node Associations tab.
   The Node Association screen opens.
3. From the **Choose Monitor** list, select the monitor you created to monitor links.
4. Click the Add button (>>) to move the monitor into the **Monitor Rule** box.
5. In the node list table, in the Associate Current Monitor Rule column, check the check box for each node address that is in the default gateway pool.
6. Click the **Apply** button.

Monitoring link performance

After you complete the Link Controller configuration, you can monitor the performance of the links by using one or more of the following tools in the Configuration utility:

- Internet Link Evaluator
- Link Statistics screens
- Link Report screens

You can use the screens to analyze the traffic patterns in your network so that you can adjust the Link Controller configuration to best meet your link management objectives. The following sections describe the screens and how to view them.

Working with the Internet Link Evaluator

The Internet Link Evaluator displays the average round trip times, the average completion rates, and the average router hops for the links in your configuration. You can use the Internet Link Evaluator to compare actual performance between links and between ISPs.

To view the Internet Link Evaluator

1. In the navigation pane, expand the **Link Statistics** item, and then click **Link Evaluator**.
   The Internet Link Evaluator screen opens.
2. For more information about interpreting the data on this screen, click the **Help** button.

◆ Note

*For additional information on the Link Evaluator, refer to Chapter 15, **Internet Link Evaluator**, in the **BIG-IP Reference Guide**.*

Working with the link statistics screens

The link statistics screens display current data for the physical and logical elements of the configuration. Each link statistics screen displays a particular aspect of your configuration.

To view the Link Statistics screens

1. In the navigation pane, expand the **Link Statistics** item, and then click one of the link statistics objects.
   The statistics screen opens for the object you selected.
2. For more information about a link statistics screen, click the **Help** button.

Working with the Link Report screen

The Link Report screen displays performance graphs for three time intervals: 30 minutes, 6 hours, and 24 hours. The graphs illustrate the volume of inbound and outbound traffic over a link during the specified time interval. The graphs also indicate any bandwidth pricing levels you have set for a link. You can view a Link Report screen for all the links in the configuration, or for a particular link in the configuration.

To view the Link Report screen for all links

1. In the navigation pane, expand the **Link Statistics** item, and then click **Links**.
   The Link Statistics screen opens.
2. Click the **Graph Link Summary** button.
   The Link Report for All Links screen opens, where you can review the bandwidth usage for all links in the most recent 30-minute, 6-hour, and 24-hour intervals.

To view the Link Report screen for a particular link

1. In the navigation pane, expand the **Link Statistics** item, and then click **Links**.
   The Link Statistics screen opens.
2. Click the **Graph Link Detail** button for the link whose data you want to review.
   The Link Report screen opens, where you can review the bandwidth usage for the particular link in the most recent 30-minute, 6-hour, and 24-hour intervals.

Additional configuration options

Whenever you configure a Link Controller, you have a number of options:

- You have the option in all configurations to configure a Link Controller redundant system for fail-over. Refer to Chapter 13, *Configuring a Redundant System*, in the ***BIG-IP Reference Guide***.
- All configurations have health monitoring options. Refer to Chapter 11, *Monitors*, in the ***BIG-IP Reference Guide***.
- When you create a pool, there is an option to set up persistence, and a choice of load balancing methods. Refer to Chapter 4, *Pools*, in the ***BIG-IP Reference Guide***.
- When you create a link, you have several advanced configuration options. Refer to Chapter 16, *Working with Link Configuration*, in the ***BIG-IP Reference Guide***.

3

Configuring Cost-Based ISP Load Balancing

- Introducing cost-based ISP load balancing
- Configuring cost-based ISP load balancing
- Additional configuration options

Introducing cost-based ISP load balancing

You can configure the Link Controller to load balance traffic based on the costs associated with traffic on each link. This configuration provides high availability and helps you control how links are used, based on the pricing structure of each link's bandwidth.

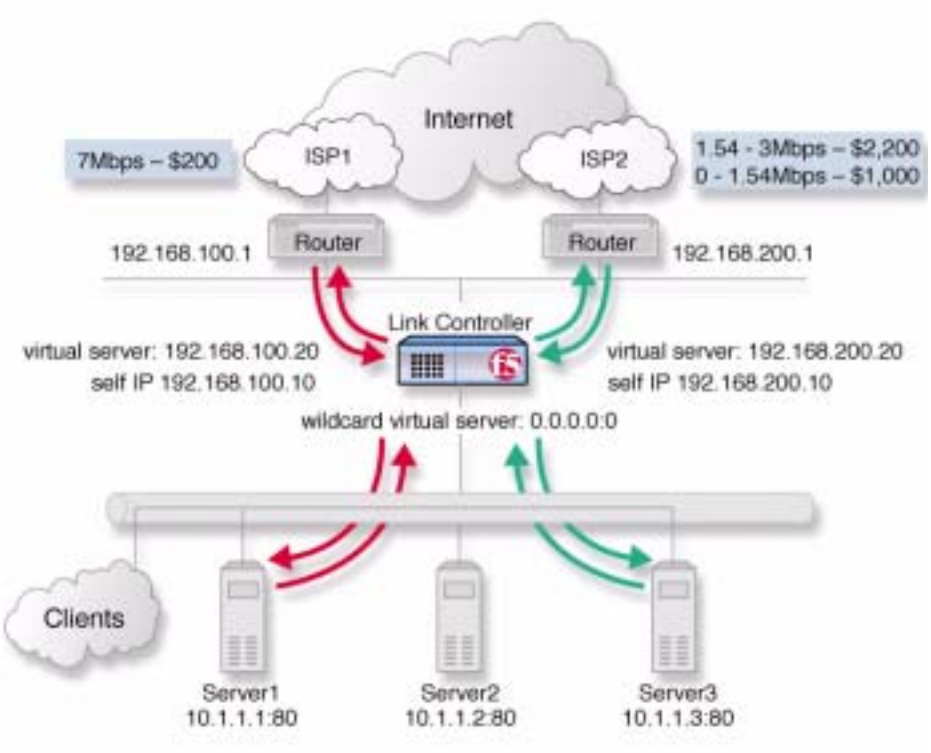

Figure 3.1 *An example of cost-based ISP load balancing*

In the example in Figure 3.1, **ISP1** handles up to .7 Mbps for $200/Mbps. The second ISP, **ISP2**, has two pricing tiers. From 0 to 1.54 Mbps, traffic costs $1000/Mbps. From 1.54 to 3 Mbps, traffic costs $2200/Mbps. In this scenario, the Link Controller directs traffic to **ISP1** if **ISP2** is unavailable, or if traffic spikes above the 3 Mbps level.

Another way to configure the Link Controller in this situation would be to direct traffic to **ISP1** if the traffic on **ISP2** exceeds 1.54 Mbps. To do this, you could set a cost limit on the **ISP2** link.

◆ Note

This type of configuration assumes that you have completed the base configuration created by the Setup utility. For more information, see Chapter 2, ***Using the Setup Utility****, in the* ***BIG-IP Reference Guide****. To use this configuration, you must configure at least three VLANs when you create the initial configuration: one VLAN for each ISP, and one VLAN for the*

internal network. The IP addresses used in this example are for demonstration only. You should substitute IP addresses appropriate for your network.

Configuring cost-based ISP load balancing

When you set up cost-based ISP, or link, load balancing, you have several tasks to complete on the Link Controller:

- **Configure the links**
  Complete the following tasks to configure the links.
  - Verify that the default gateway pool that contains the IP address of each ISP, or link, is configured correctly.
  - Add the links to the configuration.
  - Create transparent monitors to verify that the path to or through ISP is available.
- **Create two load balancing pools**
  You must define one pool that load balances the content servers. The other pool, the **default_gateway_pool**, is created when you add the IP addresses of the links while setting the default gateways in the Setup utility.
- **Configure virtual servers**
  You need to configure a virtual server in the network of the link for each ISP to load balance inbound connections across the servers. You also configure one wildcard virtual server (**0.0.0.0:0**) to load balance outbound connections across the routers.
- **Add a wide IP for inbound load balancing**
  Add a wide IP to handle inbound DNS requests for each pair of virtual servers you add for each link.
- **Configure the price weighting for each link**
  Configure the price weighting you want to associate with each link. The Link Controller uses the price weighting to direct traffic to a particular link.
- **Manage links**
  Additional configuration options are available for each link.

Defining the pools for an additional Internet connection

First, define one pool that load balances the content servers, and one pool to load balance the routers. Figure 3.2 is an example of how the network devices and servers are grouped into pools.

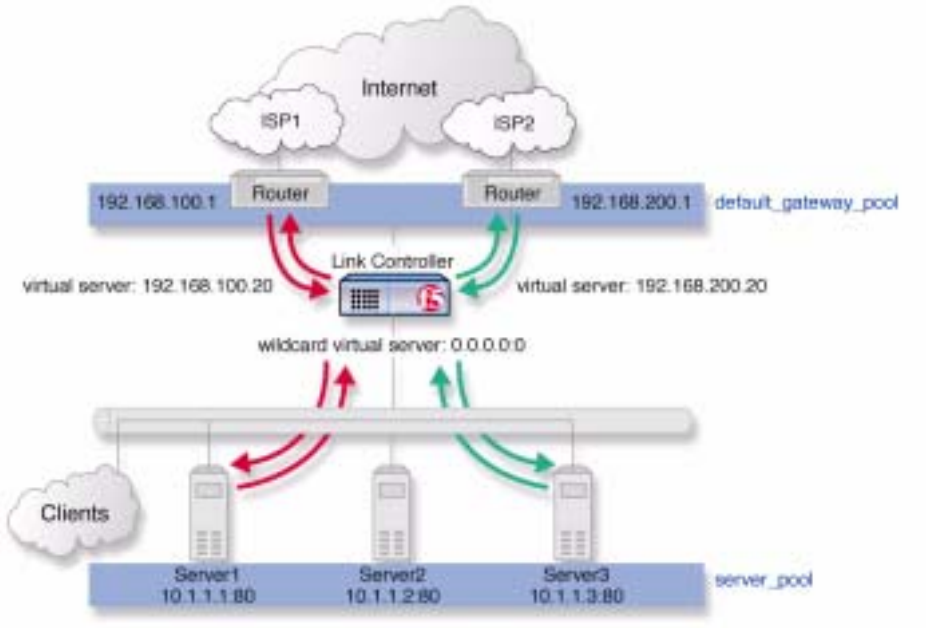

Figure 3.2 *The pools required for link load balancing*

To create the inbound load balancing pool using the Configuration utility

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

   Configuration notes

   For the example in Figure 3.1:

   Create the pool ***server_pool*** *containing the members* ***10.1.1.1:80*** ***10.1.1.2:80****, and* ***10.1.1.3:80****.*

To create the default gateway pool using the Configuration utility

If you configured more than one default gateway in the Setup utility, the Link Controller already created a **default_gateway_pool** pool. You can skip this step. If you do not have a **default_gateway_pool**, create one by completing the following task.

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

   Configuration notes

   For the example in Figure 3.1:

   *Create the pool **default_gateway_pool** containing the router inside addresses **192.168.100.1:0** and **192.168.200.1:0**.*

Defining the virtual servers for an additional Internet connection

After you create the pools, you configure the virtual servers, one for each link that load balances inbound connections across the servers. You also configure one wildcard virtual server to load balance outbound connections across the routers. Each of the virtual servers you create references either the **default_gateway_pool** or the **server_pool** in the configuration.

To define the virtual servers for inbound traffic using the Configuration utility

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

   Configuration notes

   For the example in Figure 3.1:

   *Note that you must create a virtual server for each link. For this example, create the virtual servers **192.168.100.20:80** and **192.168.200.20:80**, and use the pool, **server_pool**.*

To define a wildcard virtual server for outbound traffic using the Configuration utility

1. In the navigation pane, click **Virtual Servers**.
   The Virtual Servers Screen opens.
2. Click the **Add** button.
   The Add Virtual Server screen opens.
3. For a wildcard virtual server, use the following settings:
   - For the virtual server address, type **0.0.0.0**.
   - For the service, type **0**, or select **any** from the list.
   - For the VLAN, select **All**.
   - For the pool, select **default_gateway_pool**.
4. Click **Done**.
   Click **Help** for details on the settings on the Add Virtual Server screens.

Setting the default gateway pool

If a Link Controller does not have a predefined route for network traffic, the unit automatically sends traffic to the pool that you define as the default gateway pool. Think of the default gateway pool as a pool of default routes. In a Link Controller configuration, the default gateway pool must contain two or more gateway IP addresses, or links. If a gateway in the default gateway pool becomes inactive, existing connections through the inactive gateway are routed through another gateway in the default gateway pool.

◆ Note

*If you configure more than one default gateway in the Setup utility, the Link Controller automatically creates a **default_gateway_pool** pool.*

To set the default gateway pool using the Configuration utility

1. In the navigation pane, click **System**.
   The System screen opens.
2. Click the Properties tab.
   The Properties screen opens.
3. From the **Default Gateway Pool** list, select the pool that contains the internal IP addresses of the gateway routers.
   In this chapter's example, this is the **default_gateway_pool** pool.
4. Click **Apply**.

◆ WARNING

Default gateway IP addresses must have a corresponding self IP address/netmask combination.

Using SNAT automap for outbound traffic

Secure network address translation (SNAT) automap is automatically configured for outbound traffic so that clients receive replies through the same ISP that their requests originated from. Figure 3.3 is an example of the SNAT automap configuration for link load balancing.

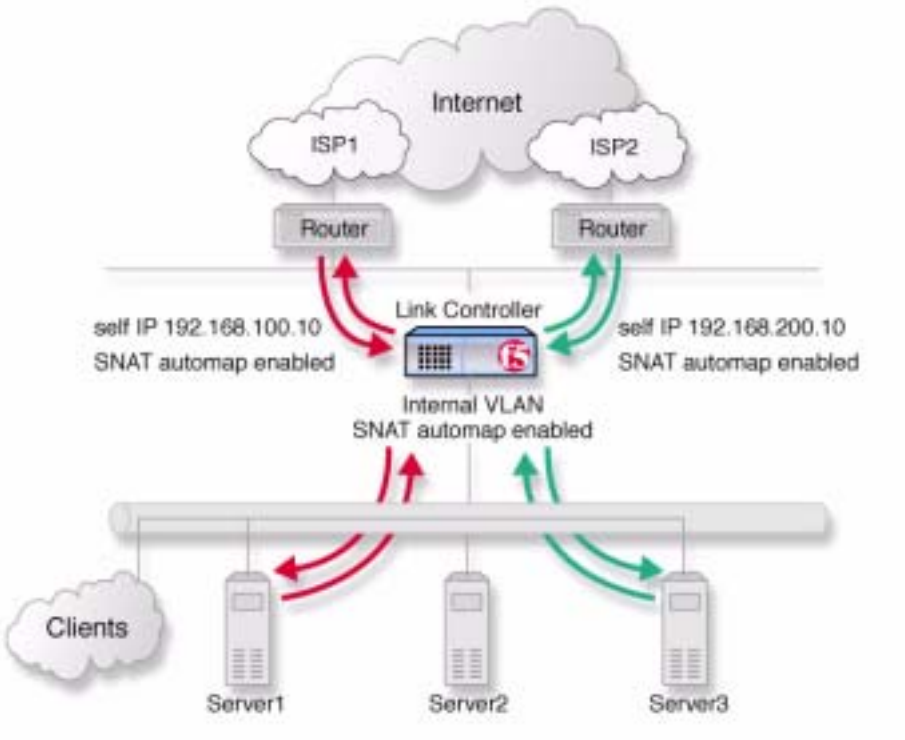

Figure 3.3 *SNAT automap feature on self IP addresses and on internal VLANs*

Adding a wide IP for inbound load balancing

To complete the link load balancing configuration, you must configure a wide IP and wide IP pool for each pair of virtual servers you created for each link. Each wide IP in your configuration has a pool of virtual servers that the Link Controller load balances incoming DNS requests to. Note that the wide IP pool is made up of only virtual servers managed by the Link Controller. When you configure the wide IP pool, you specify the load balancing methods that the Link Controller applies to incoming DNS requests.

To add a new wide IP in the Configuration utility

1. In the navigation pane, click **Link Configuration**, and then click **Inbound LB**.
   The Wide IPs List screen opens.
2. Click the **Add** button.
   The Define Wide IP (Step 1 of 2) screen opens, where you add the basic properties for the wide IP.
3. In the **Wide IP Name** box, type a name for the wide IP.
4. In the **Wide IP Port** box, type a port number to associate with the wide IP. Alternately, you can select a service, or type **0** to indicate any port.
5. Click **Next**.
   The Define Wide IP (Step 2 of 2) screen opens, where you add the virtual servers to the wide IP.
6. In the Available list, click the virtual servers that you want to add to this pool, and click the Add (-->>) button. To select more than one virtual server at a time, hold down the CTRL key on your keyboard. Note that you should include a virtual server for each link in the pool.
7. Click **Finish**.
   The Wide IP List screen opens, and you see the newly-created wide IP in the list.

Configuring price weighting for each link

After you configure the wide IP, you can set the price weighting that you want the Link Controller to use when load balancing traffic for the links, based on the billing structure that your ISP uses. The Link Controller load balances traffic to another link if the lowest cost link reaches a threshold that you have set. This helps you control traffic based on the cost the ISP is charging for the bandwidth. The following tasks describe how to configure the cost values for the scenario shown in Figure 3.1, on page 3-1.

To configure the price weighting for a link

1. In the navigation pane, expand the **Link Configuration** item, and then click **Links**.
   The Link List screen opens.
2. Click the name of one of the links in the link list.
   The Link Properties screen for that link opens.
3. Click the Link Weighting tab.
   The Link Weighting screen opens.
4. Click the **Use Price Weighting** option.
5. Next, you configure the following cost elements associated with the link:

- In the **Prepaid Segment** box, you can type the maximum bandwidth usage you pay for each month, regardless of how much you use. In the example in this document, **ISP2** charges $1000 for bandwidth usage in the 0 to 1.54 Mbps range. That means that you pay every month for up to 1.54 Mbps even if you do not use the link at all. So, using the example, you would type **1540** in the **Prepaid Segment, Up to Kbps** box.
- In the **Incremental Segment** box, type the bandwidth (in Kbps) and the associated cost of the next pricing tier. Based on the example, **ISP2** charges $2200 for bandwidth usage in the 1.54 to 3 Mbps range. Using this example, type **3000** in the **Up to Kbps** box and **2200** in the **Cost** box. Click the Add button (>>) to add the new cost tier to the configuration. You can add additional cost tiers to the configuration if required.

6. After you complete the configuration, click the **Apply** button.

Important

If you configure price weighting for one link on the Link Controller, you must configure price weighting for all of the remaining links in the configuration. If you do not, the Link Controller load balances only to the link for which price weighting is defined.

Configuring link capacity limits for load balancing

In addition to setting cost values for the bandwidth usage on a link, you can set link capacity limits for the actual traffic on the link. When you set limits on the link bandwidth capacity, you can set independent thresholds for inbound, outbound, and concurrent total traffic. There are several configuration details to consider:

- You can moderate the volume of bandwidth used for outbound requests, that is, the traffic generated by users inside the firewall, by setting a limit on outbound traffic.
- If you purchase bandwidth based on tiered pricing, you may want to limit the total traffic to a data transfer rate that keeps the volume of bandwidth used at or below a certain level of the pricing tier.
- By setting link capacity limits, you can ensure that a link does not become completely saturated before the Link Controller marks the link as **unavailable** for new traffic.

To set link capacity limits using the Configuration utility

1. In the navigation pane, expand the **Link Configuration** item, and then click **Links**.
   The Link List screen opens.
2. Click the name of one of the links in the link list.
   The Link Properties screen for that link opens.
3. Add any limit settings that you want to configure.

4. Click **Apply** to add your changes to the configuration.
   For details on the specific settings on the Link Properties screen, click the **Help** button.

Configuring transparent monitors for the links

When you create the default gateway pool, the Link Controller automatically creates simple ICMP monitors that check to make sure the IP addresses in the default gateway pool are available to the Link Controller. In addition to the default ICMP monitors, you can configure transparent monitors that verify the path taken by traffic through each link. You can use transparent monitors to check the availability of a device in an ISP network or on the Internet.

To configure a transparent monitor using the Configuration utility

1. In the navigation pane, expand the **Link Configuration** item, and then click **Links**.
   The Link List screen opens.
2. Click the name of the link that you want to modify.
   The Link Properties screen opens.
3. Click the Link Monitor tab.
   The Monitor screen opens.
4. Type in the name of your monitor, for example **LinkMonitor** (it must be different from the monitor template name), and select the **tcp_echo** monitor template.
5. Click the **Next** button.
   The Configure Basic Properties screen opens. The default **interval** and **timeout** values should be sufficient.
6. Check the **Transparent** box, and click the **Next** button.
   The Configure Destination IP and Service (Alias) screen opens.
7. In the **Destination IP** box, type the IP address of a network device on the other side of the router that you want to monitor.
   - If you are monitoring a device at the ISP, you should contact the ISP for more information about how to configure this monitor.
   - If you are monitoring a device in an ISP network, you must configure a separate monitor, with a specific IP and port combination for the device, for each ISP network.
   - If you are monitoring a device on the public Internet, you can create one monitor that is applicable for all links.

8. In the **Destination Service** box, type the service number you want to monitor. For example, if Telnet is enabled on the destination device, type **23**.
9. When you have finished configuring the monitor, click **Apply**.

Note

You can also set up ICMP transparent monitors. When your ISP does not allow TCP monitoring, use ICMP monitoring instead.

To associate the monitor with the members of the default gateway pool

After you create the monitors to check the availability of the links, you must associate the monitors with the routers in the default gateway pool.

1. In the navigation pane, click **Monitors**.
   The Monitors screen opens.
2. Click the Node Associations tab.
   The Node Association screen opens.
3. From the **Choose Monitor** list, select the monitor you created to monitor links.
4. Click the Add button (>>) to move the monitor into the **Monitor Rule** box.
5. In the node list table, in the Associate Current Monitor Rule column, check the box for each node address that is in the default gateway pool.
6. Click the **Apply** button.

Monitoring link performance

After you complete the Link Controller configuration, you can monitor the performance of the links by using one or more of the following tools in the Configuration utility:

- Internet Link Evaluator
- Link Statistics screens
- Link Report screens

You can use the screens to analyze the traffic patterns in your network so that you can adjust the Link Controller configuration to best meet your link management objectives. The following sections describe the screens and how to view them.

Working with the Internet Link Evaluator

The Internet Link Evaluator displays the average round trip times, the average completion rates, and the average router hops for the links in your configuration. You can use the Internet Link Evaluator to compare actual performance between links and between ISPs.

To view the Internet Link Evaluator

1. In the navigation pane, expand the **Link Statistics** item, and then click **Link Evaluator**.
   The Internet Link Evaluator screen opens.
2. For more information about interpreting the data on this screen, click the **Help** button.

◆ Note

*For additional information on the Link Evaluator, refer to Chapter 15, **Internet Link Evaluator**, in the **BIG-IP Reference Guide**.*

Working with the link statistics screens

The link statistics screens display current data for the physical and logical elements of the configuration. Each link statistics screen displays a particular aspect of your configuration.

To view the Link Statistics screens

1. In the navigation pane, expand the **Link Statistics** item, and then click one of the link statistics objects.
   The statistics screen for the object you selected opens.
2. For more information about a link statistics screen, click the **Help** button.

Working with the Link Report screen

The Link Report screen displays performance graphs for three time intervals: 30 minutes, 6 hours, and 24 hours. The graphs illustrate the volume of inbound and outbound traffic over a link during the specified time interval. The graphs also indicate any bandwidth pricing levels you have set for a link. You can view a Link Report screen for all the links in the configuration, or for a particular link in the configuration.

To view the Link Report screen for all links

1. In the navigation pane, expand the **Link Statistics** item, and then click **Links**.
   The Link Statistics screen opens.

2. Click the **Graph Link Summary** button.
   The Link Report for All Links screen opens, where you can review the bandwidth usage for all links in the most recent 30-minute, 6-hour, and 24-hour intervals.

To view the Link Report screen for a particular link

1. In the navigation pane, expand the **Link Statistics** item, and then click **Links**.
   The Link Statistics screen opens.
2. Click the **Graph Link Detail** button for the link whose data you want to review.
   The Link Report screen opens, where you can review the bandwidth usage for the particular link in the most recent 30-minute, 6-hour, and 24-hour intervals.

Additional configuration options

Whenever you configure a Link Controller, you have a number of options:

- You have the option in all configurations to configure a Link Controller redundant system for fail-over. Refer to Chapter 13, *Configuring a Redundant System*, in the ***BIG-IP Reference Guide***.
- All configurations have health monitoring options. Refer to Chapter 11, *Monitors*, in the ***BIG-IP Reference Guide***.
- When you create a pool, there is an option to set up persistence, and a choice of load balancing methods. Refer to Chapter 4, *Pools*, in the ***BIG-IP Reference Guide***.
- When you create a link, you have several advanced configuration options. Refer to Chapter 16, *Working with Link Configuration*, in the ***BIG-IP Reference Guide***.

4

Working with Link Controllers in a 3-DNS Sync Group

- Overview of a sync group
- Configuring a sync group

Overview of a sync group

A ***sync group*** is a group of 3-DNS Controllers and Link Controllers that share configuration information. In a sync group, the ***principal*** controller is the 3-DNS Controller that initiates metrics collection by the **big3d** agents, auto-discovers objects in the network, and is the preferred system on which to make configuration changes for the sync group. Both the principal 3-DNS Controller and the ***receiver*** controllers in the sync group receive broadcasts of metrics data from the **big3d** agents. All members of the sync group also receive broadcasts of updated configuration settings from the 3-DNS Controller that has the latest configuration changes.

Understanding how a sync group works

The sync group feature synchronizes individual configuration files, such as **wideip.conf**, and other files that store system settings. The 3-DNS Controllers and Link Controllers in a sync group operate as peer servers. At set intervals, the synchronization process compares the time stamps of the configuration files earmarked for synchronization on all of the sync group members. If the time stamp on a specific file differs between the members, the member with the latest file broadcasts the file to all of the other members in the sync group.

Planning a sync group

When you define the sync group, you select the sync group members from the list of 3-DNS Controllers and Link Controllers you have already defined in the 3-DNS configuration. In the Configuration utility, the sync group lists the controllers in the order in which you selected them. The first 3-DNS Controller in the list becomes the principal 3-DNS Controller. The remaining 3-DNS Controllers and Link Controllers in the list become receivers. If the principal 3-DNS Controller becomes disabled, the next 3-DNS Controller in the list becomes the principal 3-DNS Controller until the original principal 3-DNS Controller comes back online.

Important

In a sync group, Link Controllers can be only receiver members, and you cannot create a sync group that contains only Link Controllers. You must have at least one 3-DNS Controller in a sync group.

Configuring a sync group

The following sections describe the procedures you follow to add a Link Controller into a sync group that already has at least one 3-DNS Controller configured and working properly. You may also want to review the information in the ***3-DNS Administrator Guide***, Chapter 10, *Adding a 3-DNS Controller to an Existing Network*.

Preparing to add a Link Controller to a sync group

Before you can add a Link Controller to a sync group, you should complete the following tasks:

- Physically install the Link Controller in its data center. (For more information on hardware installation, refer to the ***Platform Guide*** that shipped with the unit.)
- Run the Setup utility on the Link Controller. (For more information on the Setup utility, see the ***BIG-IP Reference Guide***, Chapter 2, *Using the Setup Utility*.)

When you have finished this part of the setup for the Link Controller, do not make any other changes to the configuration.

Adding a Link Controller to a 3-DNS sync group

Once you have installed the Link Controller and run the Setup utility, you are ready to add the Link Controller to the sync group. There are three tasks to adding a Link Controller to a 3-DNS sync group:

- Run the **merge_configs** script on the sync group's principal 3-DNS Controller.
- Add the Link Controller to the sync group using the principal 3-DNS Controller's Configuration utility.
- Run the **3dns_add** script on the Link Controller.

The following sections explain the specific steps for each of the previous tasks. You must perform these tasks in the order they are listed.

Important

Before you add the Link Controller to the 3-DNS sync group, we recommend that you back up both the 3-DNS Controller configuration and the Link Controller configuration.

To run the merge_configs script

From the command line on the principal 3-DNS Controller, run the **merge_configs** script by typing the following command, where **<ip_address>** is the IP address of the Link Controller that you want to add to the sync group.

```
/usr/local/bin/merge_configs -peer <ip_address>
```

To make the sync group aware of the Link Controller

Using the Configuration utility on the principal 3-DNS Controller, add the Link Controller to the sync group.

1. In the navigation pane, click **3-DNS Sync**.
   The Synchronization screen opens.
2. On the toolbar, click **Add to Group**.
   The Add a 3-DNS to a Sync Group screen opens.
3. Check the box next to the Link Controller that you want to add to the sync group, and click **Add**.

To add the Link Controller to the sync group and start synchronization

The final step in adding the Link Controller to a 3-DNS sync group is to run the **3dns_add** script on the Link Controller. The script moves the synchronized configuration to the Link Controller, and finalizes the sync group setup. You can run the **3dns_add** script on the Link Controller either by using a remote secure shell session, or by using a monitor and keyboard connected directly to the controller.

1. At the **login** prompt on the new controller, type **root**.
2. At the **password** prompt, type the password you configured when you ran the Setup utility.
3. To run the script, type **3dns_add** at the command line.
   The script performs the following tasks:
   - Copies the existing controller's configuration to the new controller
   - Sets up SSH communications between the new controller and existing F5 devices in the network
   - Copies the existing controller's iQuery key to the new controller so communications between the controller and the **big3d** agents are secure

Verifying the configuration

Once the script finishes, we recommend that you verify the configuration from the principal controller in the sync group. For details on verifying the configuration, see *Verifying the configuration*, in the ***3-DNS Administrator Guide***, Chapter 10, *Adding a 3-DNS Controller to an Existing Network*.

Glossary

A record

The **A** record is the ADDRESS resource record that a Link Controller returns to a local DNS server in response to a name resolution request. The **A** record contains a variety of information, including one or more IP addresses that resolve to the requested domain name.

active unit

In a redundant system, an active unit is a system that currently load balances name resolution requests. If the active unit in the redundant system fails, the standby unit assumes control and begins to load balance requests.

alternate method

The alternate method specifies the load balancing mode that the Link Controller uses to pick a virtual server if the preferred method fails. See also *preferred method*.

Any IP Traffic

Any IP Traffic is a feature that allows the Link Controller to load balance protocols other than TCP and UDP.

big3d agent

The **big3d** agent is a monitoring agent that collects metrics information about server performance and network paths between a Link Controller and a specific local DNS server. The Link Controller uses the information collected by the **big3d** agent for dynamic load balancing.

bigpipe

The **bigpipe** utility provides command line access to the Link Controller.

BIG/stat

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

BIG/top

BIG/top is a statistical monitoring utility that ships on the Link Controller. 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**.

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.

completion rate

The completion rate is the percentage of packets that a server successfully returns during a given session.

Completion Rate mode

The Completion Rate mode is a dynamic load balancing mode that distributes connections based on which network path drops the fewest packets, or allows the fewest number of packets to time out.

default VLANs

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

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.

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 load balancing

Dynamic load balancing modes use current performance information from each node to determine which node should receive each new connection. The different dynamic load balancing modes incorporate different performance factors such as current server performance and current connection load.

dynamic load balancing modes

Dynamic load balancing modes base the distribution of name resolution requests to virtual servers on live data, such as current server performance and current connection load.

Dynamic Ratio load balancing mode

Dynamic Ratio mode is like Ratio mode (see *Ratio mode*), except that ratio weights are based on continuous monitoring of the links and are, therefore, continually changing.

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 Link Controller. 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.

fail-over

Fail-over 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.

fail-over cable

The fail-over cable directly connects the two controller units together in a redundant system.

Fastest mode

Fastest mode is a dynamic load balancing mode that bases connection distribution on which server currently exhibits the fastest response time to node pings.

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 Link Controller redundant system.

Global Availability mode

Global Availability is a static load balancing mode that bases connection distribution on a particular server order, always sending a connection to the first available server in the list. This mode differs from Round Robin mode in that it searches for an available server always starting with the first server in the list, while Round Robin mode searches for an available server starting with the next server in the list (with respect to the server selected for the previous connection request).

health check

A health check is a Link Controller feature that determines whether a node is **up** or **down**. Health checks are implemented through health monitors. See also *health monito*r 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* and *external monitor*.

host

A host is a network server that manages one or more virtual servers that the Link Controller uses for load balancing.

ICMP (Internet Control Message Protocol)

ICMP is an Internet communications protocol used to determine information about routes to destination addresses, such as virtual servers managed by Link Controller systems.

interface

An interface is the physical port on a Link Controller. See also *link*.

internal VLAN

The internal VLAN is a default VLAN on the Link Controller. 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 (Internet Security Protocol)

IPSEC is a communications protocol that provides security for the network layer of the Internet without imposing requirements on applications running above it.

iQuery

The iQuery protocol is used to exchange information between Link Controllers. The iQuery protocol is officially registered with IANA for port 4353, and works on UDP and TCP connections.

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.

last hop

A last hop is the final hop a connection took to get to the Link Controller. You can allow the Link Controller 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.

Least Connections mode

Least Connections mode is a dynamic load balancing mode that bases connection distribution on which server currently manages the fewest open connections.

link

A link is a physical interface on the Link Controller connected to another physical interface in a network.

link aggregation

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

Link Controller active unit

In a redundant system, the Link Controller active unit is the controller 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.

Link Controller web server

The Link Controller web server runs on a Link Controller and hosts the Configuration utility.

load balancing mode

A load balancing mode is a particular method of determining how to distribute connections across links.

local DNS

A local DNS is a server that makes name resolution requests on behalf of a client. With respect to the Link Controller, local DNS servers are the source of name resolution requests. Local DNS is also referred to as LDNS.

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 pool. Pools typically include multiple member nodes.

metrics information

Metrics information is the data that is typically collected about the paths between Link Controller systems, EDGE-FX Caches or GLOBAL-SITE systems, and local DNS servers. Metrics information is also collected about the performance and availability of virtual servers. Metrics information is used for load balancing, and it can include statistics such as round trip time, packet rate, and packet loss.

MindTerm SSH

MindTerm SSH is the third-party application on Link Controller systems that uses SSH for secure remote communications. SSH encrypts all network traffic (including passwords) to effectively eliminate eavesdropping, connection hijacking, and other network-level attacks. SSH also provides secure tunneling capabilities and a variety of authentication methods.

minimum active members

The minimum active members is the number of members that must be active in a priority group in order for the Link Controller 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 Link Controller 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 destination IP address or IP address:port

The monitor destination IP address or address:port for a user defined monitor is used mainly for setting up a node alias for the monitor to check. All nodes associated with that monitor will be marked down if the alias node (destination IP address:port) is marked down. See also *node alias*.

monitor instance

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

monitor template

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

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

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.

name server

A name server is a server that maintains a DNS database, and resolves domain name requests to IP addresses using that database.

named

The **named** daemon manages domain name server software.

NAT (Network Address Translation)

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

node

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

node address

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

node alias

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

node port

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

node status

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

Observed mode

Observed mode is a dynamic load balancing mode 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.

packet rate

The packet rate is the number of data packets per second processed by a server.

Packet Rate mode

The Packet Rate mode is a dynamic load balancing mode that distributes connections based on which available server currently processes the fewest packets per second.

path

A path is a logical network route between a Link Controller and a local DNS server.

path probing

Path probing is the collection of metrics data, such as round trip time and packet rate, for a given path between a requesting LDNS server and a Link Controller.

performance monitor

A performance monitor gathers statistics and checks the state of a target device.

persistence

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

picks

Picks represent the number of times a particular virtual server is selected to receive a load balanced connection.

pool

A pool is composed of a group of network devices (called members). The Link Controller 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 ratio

A pool ratio is a ratio weight applied to pools in a wide IP. If the Pool LB mode is set to Ratio, the Link Controller uses each pool for load balancing in proportion to the weight defined for the pool.

port

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

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*.

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.

Predictive mode

Predictive mode is a dynamic load balancing mode 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 mode also ranks server performance over time, and passes connections to servers which exhibit an improvement in performance rather than a decline.

preferred method

The preferred method specifies the first load balancing mode that the Link Controller uses to load balance a resolution request. See also *alternate method*.

principal 3-DNS

The principal controller is the 3-DNS Controller that initiates metrics collection by the **big3d** agents, auto-discovers objects in the network, and is the preferred system on which to make configuration changes for a sync group. All sync group members (both the principal and receivers) receive broadcasts of metrics data from the **big3d** agents throughout the network. All sync group members also receive broadcasts of updated configuration settings from the sync group member that has the latest configuration changes. See also *receiver controller, sync group*.

QOS equation

The QOS equation is the equation on which the Quality of Service load balancing mode is based. The equation calculates a score for a given path between a link and a local DNS server. The Quality of Service mode distributes connections based on the best path score for an available link. You can apply weights to the factors in the equation, such as round trip time and completion rate.

Quality of Service load balancing mode

The Quality of Service load balancing mode is a dynamic inbound load balancing mode that bases connection distribution on a configurable combination of the packet rate, completion rate, round trip time, hops, virtual server capacity, kilobytes per second, and topology information.

rate class

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

rate filter

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

ratio

A ratio is a parameter that assigns a weight to a virtual server for load balancing purposes.

Ratio mode

The Ratio load balancing mode distributes connections across an array of virtual servers in proportion to the ratio weights assigned to each individual virtual server.

receiver controller

A receiver controller is a sync group member that receives metrics data and configuration updates from the principal 3-DNS Controller. All sync group members (both the principal and receivers) receive broadcasts of metrics data from the **big3d** agents throughout the network. All sync group members also receive broadcasts of updated configuration settings from the sync group member that has the latest configuration changes. See also *principal controller, sync group*.

redundant system

Redundant system refers to a pair of controllers that are configured for fail-over. In a redundant system, there are two controller 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.

remote administrative IP address

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

resource record

A resource record is a record in a DNS database that stores data associated with domain names. A resource record typically includes a domain name, a TTL, a record type, and data specific to that record type. See also *A record.*

RFC 1918 address

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.

Round Trip Time mode

Round Trip Time mode is a dynamic load balancing mode that bases connection distribution on which virtual server has the fastest measured round trip time between the link and the local DNS server.

RTT (round trip time)

Round trip time is the calculation of the time (in microseconds) that a local DNS server takes to respond to a ping issued by the **big3d** agent running on a link. The Link Controller takes RTT values into account when it uses dynamic load balancing modes.

self IP address

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

service

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

Setup utility

The Setup utility is a utility that takes you through the initial system configuration process. The Setup utility runs automatically when you turn on a system for the first time.

SNAT (Secure Network Address Translation)

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

SNAT automap

This feature allows the Link Controller to perform a SNAT automatically on any connection that is coming from the system's internal VLAN. It is easier to use than traditional SNATs and solves certain problems associated with traditional SNATs.

SNMP (Simple Network Management Protocol)

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

source processing

Source processing means that the interface rewrites the source of an incoming packet.

SSH

SSH is a protocol for secure remote login and other secure network services over a non-secure network.

SSL gateway

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

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 Link Controller that preserves connection and persistence information in a Link Controller redundant system.

static load balancing modes

Static load balancing modes base connection distribution on a pre-defined list of criteria and virtual server availability; they do not take current server performance or current connection load into account.

sticky mask

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

STP (Spanning Tree Protocol)

STP is a protocol that provides loop resolution in configurations where one or more external switches is connected in parallel with the Link Controller.

sync group

A sync group is a group of 3-DNS Controllers and Link Controllers that synchronize system configurations and zone files (if applicable). Sync groups have one principal controller, and may contain one or more receiver controllers. All sync group members (both the principal and receivers) receive broadcasts of metrics data from the **big3d** agents throughout the network. All sync group members also receive broadcasts of updated configuration settings from the 3-DNS Controller that has the latest configuration changes. See also *principal controller*, *receiver controller*.

tagged VLAN

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

transparent node

A transparent node appears as a router to other network devices, including the Link Controller.

trunk

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

unavailable

The **unavailable** status is used for links and virtual servers. When a link or virtual server is **unavailable**, the Link Controller does not use it for load balancing.

unknown

The **unknown** status is used for links and virtual servers. When a link or virtual server is new to the Link Controller and does not yet have metrics information, the Link Controller marks its status as **unknown**. The Link Controller can use unknown servers for load balancing, but if the load balancing mode is dynamic, the Link Controller uses default metrics information for the unknown server until it receives live metrics data.

up

The **up** status is used for links and virtual servers. When a link or virtual server is **up**, the link or virtual server is available to respond to process connections.

user-defined monitor

A user-defined monitor is a custom monitor configured by a user, based on a system-supplied monitor template. For some monitor types, you must create a user-defined monitor in order to use them. For all monitor types, you must create a user-defined monitor to change the default values for system-supplied monitors.

virtual address

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

virtual port

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

virtual server

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

VLAN

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

VLAN name

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

watchdog timer card

A watchdog timer card is a hardware device that monitors the Link Controller for hardware failure.

wide IP

A wide IP is a collection of one or more domain names that maps to one or more groups of virtual servers managed either by Link Controller systems, EDGE-FX Caches, or by host servers. The Link Controller load balances name resolution requests across the virtual servers that are defined in the wide IP that is associated with the requested domain name.

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.

Index

T

U

V

W