

BIG-IP[®] Link Controller[™]: Implementations

Version 11.6



Table of Contents

Legal Notices.....	5
Acknowledgments.....	7
Chapter 1: Configuring the Link Controller system to manage traffic.....	11
Overview: Configuring the Link Controller system to manage traffic.....	11
About the initial setup of Link Controller.....	12
Task summary.....	12
Creating VLANs for communication between the network and links.....	13
Creating a default gateway pool.....	13
Creating a default route to the Internet	14
Creating links to define the physical connections to the Internet.....	14
Creating listeners to detect traffic coming from ISPs.....	15
Creating a load balancing pool.....	15
Creating virtual servers to load balance connections across servers.....	15
Creating a wildcard virtual server.....	16
Creating a wide IP that encompasses virtual servers.....	16
Implementation result.....	16
Chapter 2: Configuring cost-based load balancing.....	17
Overview: Configuring cost-based load balancing.....	17
Task summary.....	18
Creating the first cost-based link for load balancing.....	19
Creating the second cost-based link for load balancing.....	19
Creating a default gateway pool.....	20
Creating a default route to the Internet	20
Creating virtual servers to load balance connections across servers.....	20
Creating a wildcard virtual server.....	21
Creating a wide IP.....	21
Implementation result.....	22
Chapter 3: Configuring bandwidth load balancing.....	23
Overview: Configuring bandwidth load balancing	23
Task summary.....	24
Creating the first bandwidth link for load balancing.....	25
Creating the second bandwidth link for load balancing.....	25
Creating a default gateway pool.....	26
Creating a default route to the Internet	26
Creating virtual servers to load balance connections across servers.....	26
Creating a wildcard virtual server.....	27
Creating a wide IP.....	27

Implementation result.....27

Chapter 4: Creating an Active-Standby Link Controller Configuration.....29

Overview: Creating an Active-Standby Link Controller Configuration.....29

Link Controller prerequisite worksheet.....30

Task summary.....31

 Establishing a device trust between Link Controller devices.....31

 Specifying an IP address for config sync.....32

 Specifying an IP address for connection mirroring.....33

 Specifying IP addresses for failover communication.....33

 Creating a Sync-Failover device group.....34

 Verifying new traffic group membership.....35

 Syncing BIG-IP configuration between Link Controller devices36

 Enabling global traffic synchronization.....36

 Running the gtm_add script.....37

Implementation result.....37

Legal Notices

Publication Date

This document was published on December 8, 2015.

Publication Number

MAN-0535-00

Copyright

Copyright © 2014-2015, 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 user licenses. F5 reserves the right to change specifications at any time without notice.

Trademarks

AAM, Access Policy Manager, Advanced Client Authentication, Advanced Firewall Manager, Advanced Routing, AFM, Application Acceleration Manager, Application Security Manager, APM, ARX, AskF5, ASM, BIG-IP, BIG-IQ, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, Edge Portal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 [DESIGN], F5 Certified [DESIGN], F5 Networks, F5 SalesXchange [DESIGN], F5 Synthesis, f5 Synthesis, F5 Synthesis [DESIGN], F5 TechXchange [DESIGN], Fast Application Proxy, Fast Cache, FirePass, Global Traffic Manager, GTM, GUARDIAN, iApps, IBR, iCall, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, LineRate, LineRate Systems [DESIGN], Local Traffic Manager, LROS, LTM, Message Security Manager, MobileSafe, MSM, OneConnect, Packet Velocity, PEM, Policy Enforcement Manager, Protocol Security Manager, PSM, Real Traffic Policy Builder, SalesXchange, ScaleN, SDAC (except in Japan), SDC, Signalling Delivery Controller, Solutions for an application world, Software Designed Applications Services, SSL Acceleration, StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TechXchange, TMOS, TotALL, Traffic Management Operating System, Traffix (except Germany), Traffix [DESIGN] (except Germany), Transparent Data Reduction, UNITY, VAULT, vCMP, VE F5 [DESIGN], Versafe, Versafe [DESIGN], VIPRION, Virtual Clustered Multiprocessing, WebSafe, and ZoneRunner, are trademarks or service marks of F5 Networks, Inc., in the U.S. and other countries, and may not be used without F5's express written consent.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by one or more patents indicated at:
<http://www.f5.com/about/guidelines-policies/patents>

Link Controller Availability

This product is not currently available in the United States.

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.

RF Interference 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 has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful 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 ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Acknowledgments

This product includes software developed by Gabriel Forté.

This product includes software developed by Bill Paul.

This product includes software developed by Jonathan Stone.

This product includes software developed by Manuel Bouyer.

This product includes software developed by Paul Richards.

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

This product includes software developed by the Politecnico di Torino, and its contributors.

This product includes software developed by the Swedish Institute of Computer Science and its contributors.

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 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 the University of Vermont and State Agricultural College and Garrett A. Wollman.

This product includes software developed by Balazs Scheidler (bazsi@balabit.hu), which is protected under the GNU Public License.

Acknowledgments

This product includes software developed by Niels Mueller (nisse@lysator.liu.se), which is protected under the GNU Public License.

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).

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 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>.

This product includes software developed by Jared Minch.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

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

This product includes RRDtool software developed by Tobi Oetiker (<http://www.rrdtool.com/index.html>) and licensed under the GNU General Public License.

This product contains software licensed from Dr. Brian Gladman under the GNU General Public License (GPL).

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

This product includes Hypersonic SQL.

This product contains software developed by the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, and others.

This product includes software developed by the Internet Software Consortium.

This product includes software developed by Nominum, Inc. (<http://www.nominum.com>).

This product contains software developed by Broadcom Corporation, which is protected under the GNU Public License.

This product contains software developed by MaxMind LLC, and is protected under the GNU Lesser General Public License, as published by the Free Software Foundation.

This product includes Intel QuickAssist kernel module, library, and headers software licensed under the GNU General Public License (GPL).

This product includes software licensed from Gerald Combs (gerald@wireshark.org) under the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or any later version. Copyright ©1998 Gerald Combs.

This product includes software developed by Thomas Williams and Colin Kelley. Copyright ©1986 - 1993, 1998, 2004, 2007

Permission to use, copy, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Permission to modify the software is granted, but not the right to distribute the complete modified source code. Modifications are to be distributed as patches to the released version. Permission to distribute binaries produced by compiling modified sources is granted, provided you

1. distribute the corresponding source modifications from the released version in the form of a patch file along with the binaries,
2. add special version identification to distinguish your version in addition to the base release version number,
3. provide your name and address as the primary contact for the support of your modified version, and
4. retain our contact information in regard to use of the base software.

Permission to distribute the released version of the source code along with corresponding source modifications in the form of a patch file is granted with same provisions 2 through 4 for binary distributions. This software is provided "as is" without express or implied warranty to the extent permitted by applicable law.

This product contains software developed by Google, Inc. Copyright ©2011 Google, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This product includes software developed by Jeremy Ashkenas and DocumentCloud, and distributed under the MIT license. Copyright © 2010-2013 Jeremy Ashkenas, DocumentCloud.

This product includes gson software, distributed under the Apache License version 2.0. Copyright © 2008-2011 Google Inc.

This product includes jxrlib software, copyright ©2009 Microsoft Corp. All rights reserved. Distributed under the new BSD license.

This product includes node-uuid software, copyright © 2010-2012, Robert Kieffer, and distributed under the MIT license.

This product includes opensv software, which is distributed under the Apache 2.0 license.

This product includes owasp-jave-encoder software, copyright © 2014, Jeff Ichnowski, and distributed under the New BSD license.

Chapter 1

Configuring the Link Controller system to manage traffic

- *Overview: Configuring the Link Controller system to manage traffic*
- *Task summary*
- *Implementation result*

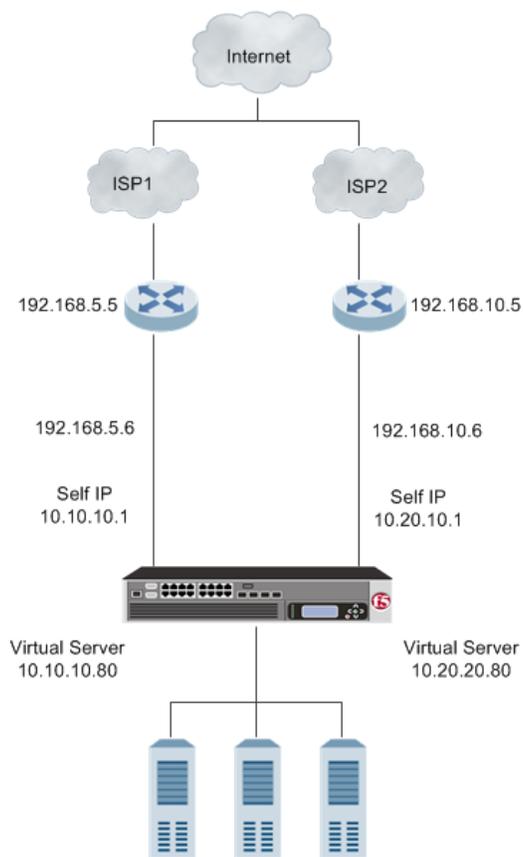
Overview: Configuring the Link Controller system to manage traffic

***Important:** This functionality is not currently available in the U.S.*

The BIG-IP® Link Controller™ provides a variety of methods for managing the traffic flowing in and out of a network. In order to integrate Link Controller into your network to manage traffic, you must complete a specific set of tasks for the capabilities of Link Controller to be available to you.

The following illustration shows a network for configuring the Link Controller system to manage Internet traffic on two connections using two different Internet service providers (ISPs).

Figure 1: Example of a network for configuring a Link Controller system to manage traffic



About the initial setup of Link Controller

Before you configure Link Controller™ on a BIG-IP® device, make sure you complete the following:

- Install the BIG-IP hardware with an initial network configuration applied.
- Provision Link Controller at the level Nominal or Dedicated.
- Configure the management IP address, network mask, and management route on the BIG-IP system.
- Designate the host name of the system as a fully qualified domain name (FQDN).
- Define the user name and password on the system that you will use when logging in to the BIG-IP Configuration utility.
- License the appropriate BIG-IP software.

Task summary

Important: This functionality is not currently available in the U.S.

Use the tasks in this implementation to integrate BIG-IP® Link Controller™ into your network. After completed, you can configure a variety of methods for managing the traffic flowing in and out of a network, including cost-based and bandwidth load balancing.

Task list

Creating VLANs for communication between the network and links

Creating a default gateway pool

Creating a default route to the Internet

Creating links to define the physical connections to the Internet

Creating listeners to detect traffic coming from ISPs

Creating a load balancing pool

Creating virtual servers to load balance connections across servers

Creating a wildcard virtual server

Creating a wide IP that encompasses virtual servers

Creating VLANs for communication between the network and links

Before you begin creating VLANs, ensure that you have completed the initial setup of the Link Controller™ system.

Create VLANs to encompass the IP addresses associated with Link Controller and the other network components that help manage DNS traffic.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. Click **Create**.
The New VLAN screen opens.
3. In the **Name** field, type the name of the first VLAN.
4. In the **Tag** field, type a numeric tag, from 1-4094, for the VLAN, or leave the field blank if you want the BIG-IP system to automatically assign a VLAN tag.
The VLAN tag identifies the traffic from hosts in the associated VLAN.
5. If you want to use Q-in-Q (double) tagging, use the **Customer Tag** setting to perform the following two steps. If you do not see the **Customer Tag** setting, your hardware platform does not support Q-in-Q tagging and you can skip this step.
 - a) From the **Customer Tag** list, select **Specify**.
 - b) Type a numeric tag, from 1-4094, for the VLAN.

The customer tag specifies the inner tag of any frame passing through the VLAN.

6. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number.
 - b) From the **Tagging** list, select **Tagged** or **Untagged**.
Select **Tagged** when you want traffic for that interface to be tagged with a VLAN ID.
 - c) If you specified a numeric value for the **Customer Tag** setting and from the **Tagging** list you selected **Tagged**, then from the **Tag Mode** list, select a value.
 - d) Click **Add**.
7. Click **Finished**.
8. Repeat these steps to create two additional VLANs.

Creating a default gateway pool

Gather the IP addresses associated with each link.

Create a default gateway pool to load balance the outbound traffic across the links.

Configuring the Link Controller system to manage traffic

1. On the Main tab, click **Local Traffic > Pools**.
The Pools list screen opens.
2. Click **Create**.
3. In the **Name** field, type a name for the pool.
Names must begin with a letter, and can contain only letters, numbers, and the underscore (_) character.

Important: *The pool name is limited to 63 characters.*

4. In the New Members area, add the IP addresses associated with each link.
 - a) In the **Address** field, type the IP address.
 - b) Click **Add**.
5. Click **Finished**.

Creating a default route to the Internet

Configure Link Controller™ to use the pool as the default gateway connection between the internal network and the Internet.

1. On the Main tab, click **Network > Routes**.
2. Click **Add**.
The New Route screen opens.
3. In the **Name** field, type a unique name.
4. From the **Resource** list, select **Use Pool**.
5. From the **Pool** list, select **default_gateway_pool**.
6. Click **Finished**.

Creating links to define the physical connections to the Internet

Gather the IP addresses of the routers associated with the ISPs and IP addresses, which correspond with the external Internet connections.

Create links using the IP addresses of routers on the network that provide a path to the Internet.

1. On the Main tab, click **Link Controller > Links**.
The Links List screen opens.
2. Click **Create**.
3. In the **Name** field, type a name for the link that represents one of the ISPs.
4. In the **Router Address** field, type the IP address of the router associated with the ISP.
5. In the **Uplink Address** field, type the IP address that corresponds with the external Internet connection.
6. In the **Service Provider** field, type the name of the ISP.
7. For the **Health Monitors** setting, from the **Available** list, select **bigip_link** and **gateway_icmp** and move the monitors to the **Enabled** list.
8. Click **Create**.
9. Repeat these steps to create the second service provider link.

Creating listeners to detect traffic coming from ISPs

Gather the self IP addresses on which Link Controller™ listens for traffic.

Create two listeners for detecting DNS traffic.

1. On the Main tab, click **Link Controller > Listeners**.
The Listeners List screen opens.
2. Click **Create**.
The Listeners properties screen opens.
3. In the **Destination** field, type the self IP address on which the Link Controller listens for traffic.
4. Click **Finished**.
5. Repeat these steps to create the second listener.

Creating a load balancing pool

Create a load balancing pool to process the inbound traffic from the Internet.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
Names must begin with a letter, and can contain only letters, numbers, and the underscore (_) character.

Important: The pool name is limited to 63 characters.

4. For the **Health Monitors** setting, from the **Available** list, select the **http** monitor and move the monitor to the **Active** list.
5. Click **Finished**.

Creating virtual servers to load balance connections across servers

Gather the IP addresses you want to use for creating the virtual servers.

Create two virtual servers, one for each link, to load balance inbound connections across the servers on the network.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. In the **Destination Address** field, type the IP address you want to use for the virtual server.
5. In the **Service Port** field, type 80, or select **HTTP** from the list.
6. Click **Finished**.
7. Repeat these steps to create the second virtual server.

Creating a wildcard virtual server

Create a wildcard server to load balance outbound connections across the routers.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For the **Destination** setting, select **Network**, and type 0.0.0.0 in the **Address** field and 0.0.0.0 in the **Mask** field.
5. In the **Service Port** field, type 0.

***Note:** Port 0 defines a wildcard virtual server that handles all types of services. If you specify a port number, you create a port-specific wildcard virtual server. In that case, the wildcard virtual server handles traffic only for the specified port.*

6. In the Resources area, from the **Default Pool** list, select **default_gateway_pool**.
7. Click **Finished**.

Creating a wide IP that encompasses virtual servers

Gather the IP addresses of the two virtual servers that you created previously.

Create a wide IP that encompasses the virtual servers.

1. On the Main tab, click **Link Controller > Inbound Wide IPs**.
The Wide IP List screen opens.
2. Click **Create**.
The New Wide IP screen opens.
3. In the **Name** field, type a name for the wide IP.
4. From the **Load Balancing Method** list, make selections from the **Preferred**, **Alternate**, and **Fallback** lists.
5. For the **Member List** field, add the virtual servers that you created previously.
 - a) From the **Virtual Server** list, select a virtual server.
 - b) Click **Add**.
6. Click **Finished**.

Implementation result

Now Link Controller™ is configured to manage the DNS traffic in and out of a network.

Chapter 2

Configuring cost-based load balancing

- *Overview: Configuring cost-based load balancing*
- *Task summary*
- *Implementation result*

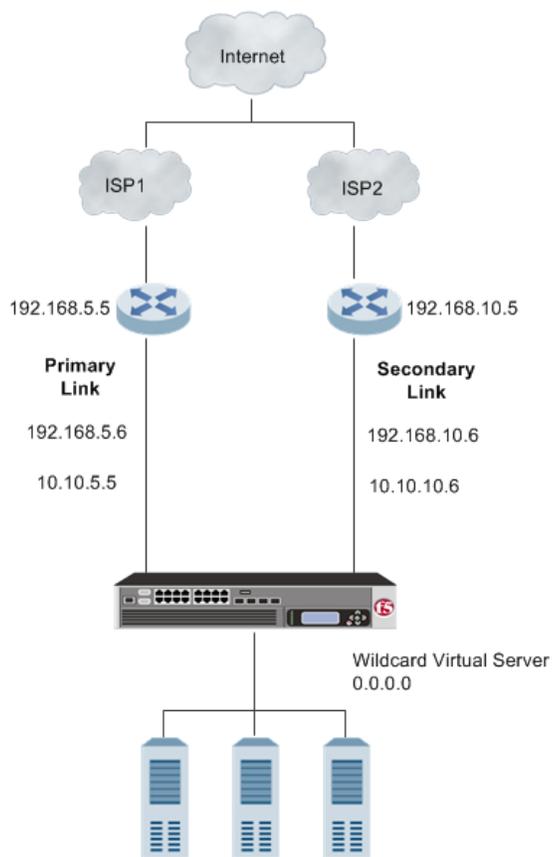
Overview: Configuring cost-based load balancing

***Important:** This functionality is not currently available in the U.S.*

You can configure BIG-IP® Link Controller™ to use cost-based load balancing to manage the traffic flowing in and out of a network. In cost-based load balancing, you prioritize link usage based on the cost of the bandwidth for the connection to the Internet. Link Controller sends traffic to the link that is currently operating at the lowest cost. As the usage cost for each link changes, Link Controller dynamically shifts traffic to the best link.

The following illustration shows a network for configuring a cost-based load balancing configuration.

Figure 2: Example of a network for configuring cost-based load balancing



As the illustration shows, while traffic remains below a certain level (for example, 4 Mbps), the Link Controller uses the Primary Link. If traffic exceeds a certain level (for example, 4 Mbps), the Link Controller sends the overflow traffic to the Secondary Link. If a link goes offline for any reason, Link Controller uses the Alternate and Fallback load balancing methods to route traffic through an available link.

Task summary

Important: This functionality is not currently available in the U.S.

Use the tasks in this implementation to create a cost-based, load balancing configuration to manage the traffic flowing in and out of a network.

Task list

- Creating the first cost-based link for load balancing*
- Creating the second cost-based link for load balancing*
- Creating a default gateway pool*
- Creating a default route to the Internet*
- Creating virtual servers to load balance connections across servers*
- Creating a wildcard virtual server*
- Creating a wide IP*

Creating the first cost-based link for load balancing

Gather the IP address of the router associated with the ISP and IP address that corresponds with the external Internet connection. In addition, determine an amount of appropriate bandwidth for the link.

Create and configure the first link on Link Controller™ to specify how traffic enters and leaves your network.

1. On the Main tab, click **Link Controller > Links**.
The Links List screen opens.
2. Click **Create**.
The New Link screen opens.
3. In the **Name** field, type a name for the link.

***Important:** Link names are limited to 63 characters.*

4. In the **Router Address** field, type the IP address of the router.
5. In the **Uplink Address** field, type the IP address that corresponds with the external Internet connection.
6. For the **Service Provider** field, type the name of the ISP provider.
7. From the **Configuration** list, select **Advanced**.
Additional controls display on the screen.
8. From the **Weighting** list, select **Price (Dynamic Ratio)**.
9. From the **Prepaid Segment** field, type the amount of bandwidth that is prepaid for the link.
10. From the **Incremental Segments** field, add the incremental segment price.
11. Click **Create**.
The Link List screen displays.

Creating the second cost-based link for load balancing

Gather the IP address of the router associated with the ISP and IP address that corresponds with the external Internet connection. In addition, determine an amount of appropriate bandwidth for the link.

Create and configure the second link on Link Controller™ to specify how traffic enters and leaves your network.

1. On the Main tab, click **Link Controller > Links**.
The Links List screen opens.
2. Click **Create**.
The New Link screen opens.
3. In the **Name** field, type a name for the link.

***Important:** Link names are limited to 63 characters.*

4. In the **Router Address** field, type the IP address of the router.
5. In the **Uplink Address** field, type the IP address that corresponds with the external Internet connection.
6. For the **Service Provider** field, type the name of the ISP provider.
7. From the **Configuration** list, select **Advanced**.
Additional controls display on the screen.
8. From the **Weighting** list, select **Price (Dynamic Ratio)**.
9. From the **Prepaid Segment** field, type the amount of bandwidth that is prepaid for the link.

10. Click **Create**.
The Link List screen displays.
11. From the **Incremental Segments** field, add the incremental segment price.

Creating a default gateway pool

Gather the IP addresses associated with each link.

Create a default gateway pool to load balance the outbound traffic across the links.

1. On the Main tab, click **Local Traffic > Pools**.
The Pools list screen opens.
2. Click **Create**.
3. In the **Name** field, type a name for the pool.
Names must begin with a letter, and can contain only letters, numbers, and the underscore (_) character.

Important: *The pool name is limited to 63 characters.*

4. In the New Members area, add the IP addresses associated with each link.
 - a) In the **Address** field, type the IP address.
 - b) Click **Add**.
5. Click **Finished**.

Creating a default route to the Internet

Configure Link Controller™ to use the pool as the default gateway connection between the internal network and the Internet.

1. On the Main tab, click **Network > Routes**.
2. Click **Add**.
The New Route screen opens.
3. In the **Name** field, type a unique name.
4. From the **Resource** list, select **Use Pool**.
5. From the **Pool** list, select **default_gateway_pool**.
6. Click **Finished**.

Creating virtual servers to load balance connections across servers

Gather the IP addresses you want to use for creating the virtual servers.

Create two virtual servers, one for each link, to load balance inbound connections across the servers on the network.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.

The New Virtual Server screen opens.

3. In the **Name** field, type a unique name for the virtual server.
4. In the **Destination Address** field, type the IP address you want to use for the virtual server.
5. In the **Service Port** field, type 80, or select **HTTP** from the list.
6. Click **Finished**.
7. Repeat these steps to create the second virtual server.

Creating a wildcard virtual server

Create a wildcard server to load balance outbound connections across the routers.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For the **Destination** setting, select **Host**, and type 0.0.0.0 in the **Address** field.
5. In the **Service Port** field, type 0.

***Note:** Port 0 defines a wildcard virtual server that handles all types of services. If you specify a port number, you create a port-specific wildcard virtual server. In that case, the wildcard virtual server handles traffic only for the specified port.*

6. Click **Finished**.

Creating a wide IP

Before you can create a wide IP, you need IP addresses from two previously created virtual servers.

Create a wide IP to which Link Controller™ load balances incoming DNS requests.

1. On the Main tab, click **Link Controller > Inbound Wide IPs**.
The Wide IP List screen opens.
2. Click **Create**.
The New Wide IP screen opens.
3. In the **Name** field, type a name for the wide IP.
4. From the **Load Balancing Method** list, make selections from the **Preferred**, **Alternate**, and **Fallback** lists.
5. For the **Member List** field, add the address of the appropriate virtual server.
 - a) From the **Virtual Server** list, select a virtual server.
 - b) Click **Add**.
6. Click **Create**.

Implementation result

You now have Link Controller™ configured to use cost-based load balancing to manage DNS traffic.

Chapter

3

Configuring bandwidth load balancing

- *Overview: Configuring bandwidth load balancing*
- *Task summary*
- *Implementation result*

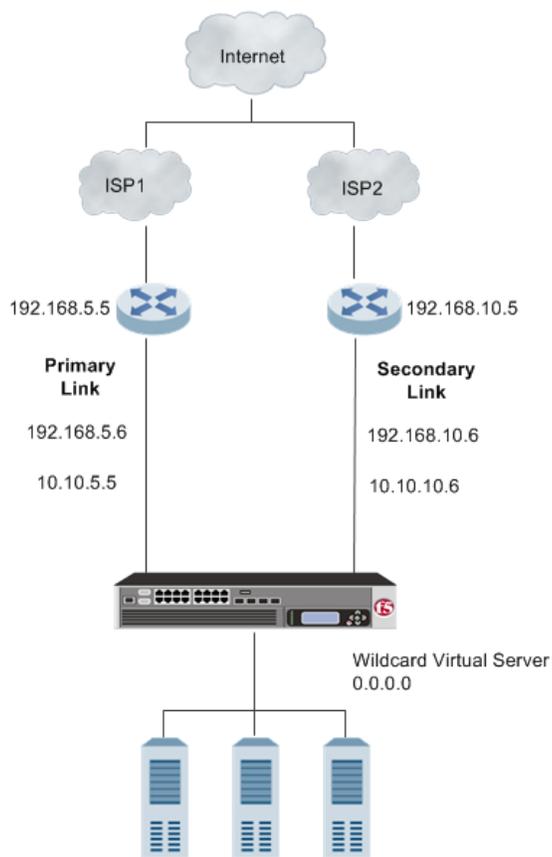
Overview: Configuring bandwidth load balancing

***Important:** This functionality is not currently available in the U.S.*

You can configure BIG-IP® Link Controller™ to use bandwidth load balancing to manage the traffic flowing in and out of a network. In bandwidth load balancing, Link Controller uses a specific link until a traffic threshold has been met. After that threshold is met, the Link Controller shifts traffic to another link. When the traffic falls below the threshold, the Link Controller shifts traffic back to the first link.

The following illustration shows a network for configuring a bandwidth load balancing configuration.

Figure 3: Example of a network for configuring bandwidth load balancing



As the illustration shows, the most cost-efficient configuration is to have the Primary Link handle traffic until it reaches a certain level (for example, 50 Mbps), then send any traffic over a certain level (for example: 50 Mbps) to the Secondary Link. When the traffic decreases, the Link Controller must switch back to using only the Primary Link again.

Task summary

Important: *This functionality is not currently available in the U.S.*

Use the tasks in this implementation to create a bandwidth load balancing configuration to manage the traffic flowing in and out of a network.

Task list

- Creating the first bandwidth link for load balancing*
- Creating the second bandwidth link for load balancing*
- Creating a default gateway pool*
- Creating a default route to the Internet*
- Creating virtual servers to load balance connections across servers*
- Creating a wildcard virtual server*
- Creating a wide IP*

Creating the first bandwidth link for load balancing

Gather the IP address of the router associated with the ISP and IP address that corresponds with the external Internet connections. For the link, determine the total bandwidth threshold you want to use.

Create and configure the first link on Link Controller™, which determines how traffic enters and leaves your network.

1. On the Main tab, click **Link Controller > Links**.
The Links List screen opens.
2. Click **Create**.
The New Link screen opens.
3. In the **Name** field, type a name for the link.

Important: Link names are limited to 63 characters.

4. In the **Router Address** field, type the IP address of the router.
5. In the **Uplink Address** field, type the IP address that corresponds with the external Internet connection.
6. For the **Service Provider** field, type the name of the ISP provider.
7. From the **Configuration** list, select **Advanced**.
Additional controls display on the screen.
8. From the **Traffic Links** field, set the total bandwidth thresholds for the link.
9. Click **Create**.

Creating the second bandwidth link for load balancing

Gather the IP address of the router associated with the ISP and IP address that corresponds with the external Internet connection. For the link, determine the total bandwidth threshold you want to use.

Create and configure the second link on Link Controller™, which determines how traffic enters and leaves your network.

1. On the Main tab, click **Link Controller > Links**.
The Links List screen opens.
2. Click **Create**.
The New Link screen opens.
3. In the **Name** field, type a name for the link.

Important: Link names are limited to 63 characters.

4. In the **Router Address** field, type the IP address of the router.
5. In the **Uplink Address** field, type the IP address that corresponds with the external Internet connection.
6. For the **Service Provider** field, type the name of the ISP provider.
7. From the **Configuration** list, select **Advanced**.
Additional controls display on the screen.
8. From the **Traffic Links** field, set the total bandwidth thresholds for the link.
9. Click **Create**.

Creating a default gateway pool

Gather the IP addresses associated with each link.

Create a default gateway pool to load balance the outbound traffic across the links.

1. On the Main tab, click **Local Traffic > Pools**.
The Pools list screen opens.
2. Click **Create**.
3. In the **Name** field, type a name for the pool.
Names must begin with a letter, and can contain only letters, numbers, and the underscore (_) character.

Important: *The pool name is limited to 63 characters.*

4. In the New Members area, add the IP addresses associated with each link.
 - a) In the **Address** field, type the IP address.
 - b) Click **Add**.
5. Click **Finished**.

Creating a default route to the Internet

Configure Link Controller™ to use the pool as the default gateway connection between the internal network and the Internet.

1. On the Main tab, click **Network > Routes**.
2. Click **Add**.
The New Route screen opens.
3. In the **Name** field, type a unique name.
4. From the **Resource** list, select **Use Pool**.
5. From the **Pool** list, select **default_gateway_pool**.
6. Click **Finished**.

Creating virtual servers to load balance connections across servers

Gather the IP addresses you want to use for creating the virtual servers.

Create two virtual servers, one for each link, to load balance inbound connections across the servers on the network.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. In the **Destination Address** field, type the IP address you want to use for the virtual server.
5. In the **Service Port** field, type 80, or select **HTTP** from the list.

6. Click **Finished**.
7. Repeat these steps to create the second virtual server.

Creating a wildcard virtual server

Create a wildcard server to load balance outbound connections across the routers.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For the **Destination** setting, select **Host**, and type 0.0.0.0 in the **Address** field.
5. In the **Service Port** field, type 0.

***Note:** Port 0 defines a wildcard virtual server that handles all types of services. If you specify a port number, you create a port-specific wildcard virtual server. In that case, the wildcard virtual server handles traffic only for the specified port.*

6. Click **Finished**.

Creating a wide IP

Before you can create a wide IP, you need IP addresses from two previously created virtual servers.

Create a wide IP to which Link Controller™ load balances incoming DNS requests.

1. On the Main tab, click **Link Controller > Inbound Wide IPs**.
The Wide IP List screen opens.
2. Click **Create**.
The New Wide IP screen opens.
3. In the **Name** field, type a name for the wide IP.
4. From the **Load Balancing Method** list, make selections from the **Preferred**, **Alternate**, and **Fallback** lists.
5. For the **Member List** field, add the address of the appropriate virtual server.
 - a) From the **Virtual Server** list, select a virtual server.
 - b) Click **Add**.
6. Click **Create**.

Implementation result

You now have Link Controller™ configured to use bandwidth load balancing to manage DNS traffic.

Chapter 4

Creating an Active-Standby Link Controller Configuration

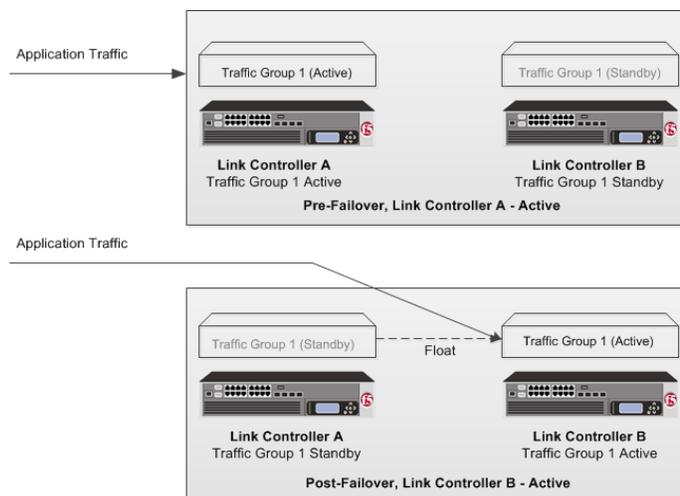
- *Overview: Creating an Active-Standby Link Controller Configuration*
- *Link Controller prerequisite worksheet*
- *Task summary*
- *Implementation result*

Overview: Creating an Active-Standby Link Controller Configuration

You can configure BIG-IP® Link Controller™ in an active-standby configuration, which is a set of two Link Controller systems: one operating as the active unit, the other operating as the standby unit. If the active unit in the active traffic group goes offline, the standby unit immediately assumes responsibility for managing traffic. The new active unit remains active until another event occurs that would cause the unit to go offline, or you manually reset the status of each unit.

This illustration shows Link Controller devices in an active-standby configuration.

Figure 4: Example of Link Controller devices in an active-standby configuration



Link Controller prerequisite worksheet

Before you set up an active-standby BIG-IP® Link Controller™ configuration, you must configure these BIG-IP components on each device that you intend to include in the device group.

Table 1: Link Controller deployment worksheet

Configuration component	Considerations
Hardware, licensing, and provisioning	Devices in a device group must match with respect to product licensing and module provisioning. Heterogeneous hardware platforms within a device group are supported.
BIG-IP software version	Each device must be running BIG-IP version 11.x. This ensures successful configuration synchronization.
Management IP addresses	Each device must have a management IP address, a network mask, and a management route defined.
FQDN	Each device must have a fully-qualified domain name (FQDN) as its host name.
User name and password	Each device must have a user name and password defined on it that you will use when logging in to the BIG-IP Configuration utility.
root folder properties	The platform properties for the root folder must be set correctly (<code>Sync-Failover</code> and <code>traffic-group-1</code>).
VLANs	You must create these VLANs on each device, if you have not already done so: <ul style="list-style-type: none"> • A VLAN for the internal network, named <code>internal</code>. • A VLAN for the external network, named <code>external</code>. • A VLAN for failover communications, named <code>HA</code>.
Self IP addresses	You must create these self IP addresses on each device, if you have not already done so: <ul style="list-style-type: none"> • Two self IP addresses (floating and non-floating) on the same subnet for VLAN <code>internal</code>. • Two self IP addresses (floating and non-floating) on the same subnet for VLAN <code>external</code>. • A non-floating self IP address on the internal subnet for VLAN <code>HA</code>. <hr/> <p><i>Note: When you create floating self IP addresses, the BIG-IP system automatically adds them to the default floating traffic group, <code>traffic-group-1</code>. To add a self IP address to a different traffic group, you must modify the value of the self IP address Traffic Group property.</i></p> <hr/> <p>Important: If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the IP address you specify must be the floating IP address for high availability fast failover that you configured for the EC2 instance.</p> <hr/>
Port lockdown	For self IP addresses that you create on each device, you should verify that the Port Lockdown setting is set to Allow All , All Default , or Allow Custom . Do not specify None .
Application-related objects	You must create any virtual IP addresses and optionally, SNAT translation addresses, as part of the local traffic configuration. You must also configure any iApp®

Configuration component	Considerations
	application services if they are required for your application. When you create these addresses or services, the objects automatically become members of the default traffic group <code>traffic-group-1</code> .
Time synchronization	The times set by the NTP service on all devices must be synchronized. This is a requirement for configuration synchronization to operate successfully.
Device certificates	Verify that each device includes an x509 device certificate. Devices with device certificates can authenticate, and thus, trust one another, which is a prerequisite for device-to-device communication and data exchange.

Task summary

Use the tasks in this implementation to create a two-member device group, with one active traffic group that syncs the BIG-IP® configuration to the peer device and provides failover capability if the peer device goes offline.

Task list

Establishing a device trust between Link Controller devices
Specifying an IP address for config sync
Specifying an IP address for connection mirroring
Specifying IP addresses for failover communication
Creating a Sync-Failover device group
Verifying new traffic group membership
Syncing BIG-IP configuration between Link Controller devices
Enabling global traffic synchronization
Running the gtm_add script

Establishing a device trust between Link Controller devices

Before you begin this task, verify that:

- Each BIG-IP® device that is to be part of the local trust domain has a device certificate installed on it.
- The local device is designated as a certificate signing authority.

Establish trust among devices on one or more network segments to constitute the *local trust domain*.

Note: *A device must be a member of the local trust domain prior to joining a device group.*

By default, the BIG-IP software includes a local trust domain with one member, which is the local device. You can choose any one of the BIG-IP devices slated for a device group and log into that device to add other devices to the local trust domain.

1. On the Main tab, click **Device Management** > **Device Trust**, and then either **Peer List** or **Subordinate List**.
2. Click **Add**.

3. Type a device IP address, administrator user name, and administrator password for the remote BIG-IP® device with which you want to establish trust. The IP address you specify depends on the type of BIG-IP device:
 - If the BIG-IP device is an appliance, type the management IP address for the device.
 - If the BIG-IP device is a VIPRION® device that is not licensed and provisioned for vCMP®, type the primary cluster management IP address for the cluster.
 - If the BIG-IP device is a VIPRION device that is licensed and provisioned for vCMP, type the cluster management IP address for the guest.
 - If the BIG-IP device is an Amazon Web Services EC2 device, type one of the Private IP addresses created for this EC2 instance.
4. Click **Retrieve Device Information**.
5. Verify that the certificate of the remote device is correct.
6. Verify that the management IP address and name of the remote device are correct.
7. Click **Finished**.

As a result of these steps, the local device is now a member of the local trust domain. Also, the BIG-IP system automatically creates a special Sync-Only device group for the purpose of synchronizing trust information among the devices in the local trust domain, on an ongoing basis.

Repeat these steps to specify each device that you want to add to the local trust domain.

Specifying an IP address for config sync

Before configuring the config sync address, verify that all devices in the device group are running the same version of BIG-IP® system software.

You perform this task to specify the IP address on the local device that other devices in the device group will use to synchronize their configuration objects to the local device.

Note: You must perform this task locally on each device in the device group.

1. Confirm that you are logged in to the device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the **Device Connectivity** menu, choose ConfigSync.
5. For the **Local Address** setting, retain the displayed IP address or select another address from the list.
F5 Networks recommends that you use the default value, which is the self IP address for the internal VLAN. This address must be a non-floating self IP address and not a management IP address.

Important: If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the internal self IP address that you select must be an internal private IP address that you configured for this EC2 instance as the **Local Address**.

6. Click **Update**.

After performing this task, the other devices in the device group can synchronize their configurations to the local device whenever a sync operation is initiated.

Specifying an IP address for connection mirroring

Specify the local self IP address that you want other devices in a device group to use when mirroring their connections to this device. Connection mirroring ensures that in-process connections for an active traffic group are not dropped when failover occurs.

Note: You must perform these steps locally on each device in the device group.

1. Confirm that you are logged in to the device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the Device Connectivity menu, choose Mirroring.
5. For the **Primary Local Mirror Address** setting, retain the displayed IP address or select another address from the list.

The recommended IP address is the self IP address for either VLAN `HA` or VLAN `internal`.

Important: If the BIG-IP device you are configuring is accessed using Amazon Web Services, then the self IP address you specify must be one of the private IP addresses that you configured for this EC2 instance as the **Primary Local Mirror Address**.

6. For the **Secondary Local Mirror Address** setting, retain the default value of **None**, or select an address from the list.
This setting is optional. The system uses the selected IP address in the event that the primary mirroring address becomes unavailable.
7. Click **Update**.

In addition to specifying an IP address for mirroring, you must also enable connection mirroring on the relevant virtual servers on this device.

Specifying IP addresses for failover communication

Specify the local IP addresses that you want other devices in the device group to use for continuous health-assessment communication with the local device. You must perform the steps locally on each device in the device group.

Note: The IP addresses that you specify must belong to route domain 0.

1. Confirm that you are logged in to the device you want to configure.
2. On the Main tab, click **Device Management > Devices**.
This displays a list of device objects discovered by the local device.
3. In the Name column, click the name of the device to which you are currently logged in.
4. From the **Device Connectivity** menu, choose Failover Network.
5. For the Failover Unicast Configuration settings, click **Add** for each IP address on this device that other devices in the device group can use to exchange failover messages with this device. The unicast IP addresses you specify depend on the type of device:

Platform	Action
Appliance without vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN _{HA}) and the static management IP address currently assigned to the device.
Appliance with vCMP	Type a static self IP address associated with an internal VLAN (preferably VLAN _{HA}) and the unique management IP address currently assigned to the guest.
VIPRION without vCMP[®]	Type a static self IP address associated with an internal VLAN (preferably VLAN _{HA}). If you choose to specify unicast addresses only (and not a multicast address), you must also type the existing, static management IP addresses that you previously configured for all slots in the cluster. If you choose to specify one or more unicast addresses and a multicast address, then you do not need to specify the existing, per-slot static management IP addresses when configuring addresses for failover communication.
VIPRION with vCMP	Type a self IP address that is defined on the guest and associated with an internal VLAN on the host (preferably VLAN _{HA}). If you choose to specify unicast failover addresses only (and not a a multicast address), you must also type the existing, virtual static management IP addresses that you previously configured for all slots in the guest's virtual cluster. If you choose to specify one or more unicast addresses and a multicast address, you do not need to specify the existing, per-slot static and virtual management IP addresses when configuring addresses for failover communication.

Important: Failover addresses should always be static, not floating, IP addresses.

6. To enable the use of a failover multicast address on a VIPRION[®] platform (recommended), then for the **Use Failover Multicast Address** setting, select the **Enabled** check box.
7. If you enabled **Use Failover Multicast Address**, either accept the default **Address** and **Port** values, or specify values appropriate for the device.
If you revise the default **Address** and **Port** values, but then decide to revert to the default values, click **Reset Defaults**.
8. Click **Update**.

After you perform these steps, other devices in the device group can send failover messages to the local device using the specified IP addresses.

Creating a Sync-Failover device group

This task establishes failover capability between two or more BIG-IP[®] devices. If an active device in a Sync-Failover device group becomes unavailable, the configuration objects fail over to another member of the device group and traffic processing is unaffected. You perform this task on any one of the authority devices within the local trust domain.

Repeat this task for each Sync-Failover device group that you want to create for your network configuration.

1. On the Main tab, click **Device Management > Device Groups**.
2. On the Device Groups list screen, click **Create**.
The New Device Group screen opens.
3. Type a name for the device group, select the device group type **Sync-Failover**, and type a description for the device group.
4. From the **Configuration** list, select **Advanced**.

5. In the Configuration area of the screen, select a host name from the **Available** list for each BIG-IP device that you want to include in the device group, including the local device. Use the Move button to move the host name to the **Includes** list.

The **Available** list shows any devices that are members of the device's local trust domain but not currently members of a Sync-Failover device group. A device can be a member of one Sync-Failover group only.

6. For the **Network Failover** setting, select or clear the check box:
 - Select the check box if you want device group members to handle failover communications by way of network connectivity. This choice is required for active-active configurations.
 - Clear the check box if you want device group members to handle failover communications by way of serial cable (hard-wired) connectivity.

For active-active configurations, you must select network failover, as opposed to serial-cable (hard-wired) connectivity.

7. For the **Automatic Sync** setting, select or clear the check box:
 - Select the check box when you want the BIG-IP system to automatically sync the BIG-IP configuration data whenever a config sync operation is required. In this case, the BIG-IP system syncs the configuration data whenever the data changes on any device in the device group.
 - Clear the check box when you want to manually initiate each config sync operation. In this case, F5 networks recommends that you perform a config sync operation whenever configuration data changes on one of the devices in the device group.

8. For the **Full Sync** setting, select or clear the check box:
 - Select the check box when you want all sync operations to be full syncs. In this case, the BIG-IP system syncs the entire set of BIG-IP configuration data whenever a config sync operation is required.
 - Clear the check box when you want all sync operations to be incremental (the default setting). In this case, the BIG-IP system syncs only the changes that are more recent than those on the target device. When you select this option, the BIG-IP system compares the configuration data on each target device with the configuration data on the source device and then syncs the delta of each target-source pair.

If you enable incremental synchronization, the BIG-IP system might occasionally perform a full sync for internal reasons. This is a rare occurrence and no user intervention is required.

9. In the **Maximum Incremental Sync Size (KB)** field, retain the default value of 1024, or type a different value.

This value specifies the total size of configuration changes that can reside in the incremental sync cache. If the total size of the configuration changes in the cache exceeds the specified value, the BIG-IP system performs a full sync whenever the next config sync operation occurs.

10. Click **Finished**.

You now have a Sync-Failover type of device group containing BIG-IP devices as members.

Verifying new traffic group membership

Verify new traffic group membership to ensure that the IP addresses in the traffic group failover. Confirm that the same traffic group (typically the default traffic group, `traffic-group-1`) has all the the appropriate floating IP addresses, such as the internal and external self IP addresses, and a virtual IP address.

Note: You must perform these steps on each Link Controller device.

1. On the Main tab, click **Device Management > Traffic Groups**.

2. In the Name column, click the name of the traffic group for which you want to view the associated objects.
This displays a list of all failover objects for the traffic group.
3. In the Address column, for the traffic group selected, verify the listing of all of the appropriate floating IP addresses, including the internal and external self IP addresses, and a virtual IP address. If a SNAT address was created, also verify it is listed.

Syncing BIG-IP configuration between Link Controller devices

Before you sync the configuration, verify that the BIG-IP® Link Controller™ devices targeted for config sync are members of a device group and that device trust is established.

Synchronize the BIG-IP configuration data from a local device to devices in a device group to ensure that all devices operate properly. When synchronizing self IP addresses, the BIG-IP system synchronizes floating self IP addresses only.

Important: You can perform these steps on either of the two devices, but not both.

1. On the Main tab, click **Device Management > Overview**.
2. In the Device Groups area of the screen, in the Name column, select the name of the relevant device group.
The screen expands to show a summary and details of the sync status of the selected device group, as well as a list of the individual devices within the device group.
3. In the Devices area of the screen, in the Sync Status column, select the device that shows a sync status of *Changes Pending*.
4. In the Sync Options area of the screen, select **Sync Device to Group**.
5. Click **Sync**.
The BIG-IP system syncs the configuration data of the selected device in the Device area of the screen to the other members of the device group.

Except for non-floating self IP addresses, the entire set of BIG-IP configuration data is replicated on each device in the device group.

Enabling global traffic synchronization

Enable global traffic synchronization options and create a name for the global traffic synchronization group.

Important: Perform the following steps only on the active system.

1. On the Main tab, click **System > Configuration > Global Traffic > General**.
The General configuration screen opens.
2. Select the **Synchronize** check box.
3. Select the **Synchronize DNS Zone Files** check box.
4. In the **Group Name** field, type the name of the synchronization group.
5. Click **Update**.

Running the gtm_add script

Before you run the `gtm_add` script, make sure that the TCP port 4353 is available on both the active and standby systems.

Run the `gtm_add` script for the standby system to acquire the configuration established on the active system.

1. On the standby system, log in to the command-line interface.
2. Type `gtm_add`, and press Enter.
3. Press the `y` key to start the `gtm_add` script.
4. Type the IP address of the active system.
5. Press Enter.
The `gtm_add` process begins, acquiring configuration data from the active system.

Implementation result

You now have created a active-standby configuration consisting of two Link Controller™ systems: one operating as the active unit, the other operating as the standby unit.

Index

A

- active-standby link controller system configuration
 - overview 29
- active-standby configuration
 - result 37

B

- bandwidth link for load balancing
 - creating 25
- bandwidth load balancing
 - overview 23
 - result 27
- bandwidth load balancing task summary
 - creating 24

C

- config sync addresses
 - specifying 32
- configuration synchronization
 - syncing to group 36
- connection mirroring
 - configuring 33
- connections
 - preserving on failover 33
- cost-based link for load balancing
 - creating 19
- cost-based load balancing task summary
 - creating 18

D

- default gateway pool
 - creating 13, 20, 26
- device discovery
 - for device trust 31
- device groups
 - creating 34
- devices
 - and mirroring limit 33

F

- failover IP addresses
 - specifying 33

G

- gtm_add script
 - running 37

I

- inbound wide IPs
 - creating 16, 21, 27

L

- Link Controller
 - setup 12
- Link Controller deployment worksheet 30
- Link Controller device trust
 - establishing 31
- Link Controller systems
 - configuring 11
 - overview 11
- Link Controller task summary
 - configuring to manage traffic 12
- links
 - creating 14
- listeners
 - creating 15
- load balancing cost-based
 - overview 17
 - result 22
- load balancing pool
 - creating 15
- local trust domain
 - and device groups 34
 - defined 31

M

- manage traffic
 - result 16

N

- network failover
 - configuring 34

P

- pool
 - creating 13, 15, 20, 26

R

- route
 - creating default 14, 20, 26

S

- Sync-Failover device groups
 - creating 34
- synchronization group
 - creating 36

T

- traffic group membership
 - verifying 35

Index

two-member device group task summary
creating 31

V

virtual servers
creating 15, 20, 26

VLANs
creating 13

W

wildcard virtual servers
creating 16, 21, 27