




*Analyze* | *Assure* | *Accelerate™*

## **Spirent Communications Step-By-Step Test Methodologies**

# **TRT Edition: High Availability Routing**

December 2004



## **Spirent Communications, Inc.**

26750 Agoura Road  
Calabasas, CA  
91302 USA

### **Support Contacts**

#### **U.S.**

E-mail: [support@spirentcom.com](mailto:support@spirentcom.com)

Web: <http://support.spirentcom.com>

Toll Free: 1-800-SPIRENT (1-800-774-7368)

Phone: + 1 818-676-2616

Fax: +1 818-880-9154

#### **Europe**

E-mail: [support.europe@spirentcom.com](mailto:support.europe@spirentcom.com)

Web: <http://support.spirentcom.com>

Phone: +33 (0) 1 61 37 22 70

Fax: +33 (0) 1 61 37 22 51

### **Copyright**

© 2004 Spirent Communications, Inc. All Rights Reserved.

All of the company names and/or brand names and/or product names referred to in this document, in particular, the name "Spirent" and its logo device, are either registered trademarks or trademarks of Spirent plc and its subsidiaries, pending registration in accordance with relevant national laws. All other registered trademarks or trademarks are the property of their respective owners. The information contained in this document is subject to change without notice and does not represent a commitment on the part of Spirent Communications. The information in this document is believed to be accurate and reliable, however, Spirent Communications assumes no responsibility or liability for any errors or inaccuracies that may appear in the document.

### **Limited Warranty**

Spirent Communications, Inc. ("Spirent") warrants that its Products will conform to the description on the face of order, that it will convey good title thereto, and that the Product will be delivered free from any lawful security interest or other lien or encumbrance.

Spirent further warrants to Customer that hardware which it supplies and the tangible media on which it supplies software will be free from significant defects in materials and workmanship for a period of twelve (12) months, except as otherwise noted, from the date of delivery (the "Warranty Period"), under normal use and conditions.

To the extent the Product is or contains software ("Software"), Spirent also warrants that the Software, if properly used by Licensee in accordance with the Software License Agreement, shall operate in material conformity with the specification for such Software for a period of ninety (90) days from the date of delivery (the "Warranty Period"). Spirent does not warrant that the functions contained in the Software will meet a specific requirement or that the operation will be uninterrupted or error free. Spirent shall have no warranty obligations whatsoever with respect to any Software which has been modified in any manner by Customer or any third party.

Defective Products and Software under warranty shall be, at Spirent's discretion, repaired or replaced or a credit issued to Customer's account for an amount equal to the price paid for such Product provided that: (a) such Product is returned to Spirent after first obtaining a return authorization number and shipping instructions, freight prepaid, to Spirent's location in the United States; (b) Customer provides a written explanation of the defect or Software failure claimed by Customer; and (c) the claimed defect actually exists and was not caused by neglect, accident, misuse, improper installation, improper repair, fire, flood, lightning, power surges, earthquake, or alteration. Spirent will ship repaired Products to Customer, freight prepaid. Except as otherwise stated, any claim on account of defective materials or for any other cause whatsoever will conclusively be deemed waived by Customer unless written notice thereof is given to Spirent within the Warranty Period. Spirent reserves the right to change the warranty and service policy set forth above at any time, after reasonable notice and without liability to Customer.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, ALL IMPLIED WARRANTIES, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, ARE HEREBY EXCLUDED, AND THE LIABILITY OF SPIRENT, IF ANY, FOR DAMAGE RELATING TO ANY ALLEGEDLY DEFECTIVE PRODUCT SHALL BE LIMITED TO THE ACTUAL PRICE PAID BY THE CUSTOMER FOR SUCH PRODUCT. THE PROVISIONS SET FORTH ABOVE STATE SPIRENT'S ENTIRE RESPONSIBILITY AND CUSTOMER'S SOLE AND EXCLUSIVE REMEDY WITH RESPECT TO ANY BREACH OF ANY WARRANTY.

# Table of Contents

## **Introduction: Protocols and Testing**

How to Limit Downtime.....	1
Truth in Testing.....	1
Using the TeraRouting Tester.....	1
Routing Protocols.....	2
Explanation of the Test Methodologies.....	2

## **The Test Methodologies**

Redundancy Protocols: VRRP/HSRP (TM-0186).....	3
BGP Graceful Restart (TM-0187).....	11
Graceful Restart for Other Protocols (TM-0188).....	21

## **Supplementary Information**

Glossary.....	28
Spirent Communications Test Methodologies Information.....	29

# Introduction: Protocols and Testing

## How to Limit Downtime

Unscheduled downtime is intolerable for any business that is financially dependent on the operation of its networks, which is why outages must be controlled through a quick system recovery.

The communications industry calls this process High Availability (HA) Routing or Non-Stop Forwarding (NSF). The fundamental principle calls for trying to minimize or eliminate any downtime. While major router vendors, service providers and the Internet Engineering Task Force (IETF) pursue solutions for reducing downtime, 100 percent uptime continues to be elusive. Instead, the proverbial “five nines” or 99.999 percent network availability is the realistic objective.

Vendors have proposed the Hot Standby Router Protocol (HSRP) and the Virtual Router Redundancy Protocol (VRRP) to reduce the recovery period. These two competing solutions describe failover processes for router components including the failure of the entire router. Other industry proposals promote circuit redundancy (load balancing) and backup links.

Yet other ideas want to minimize downtime during control plane events such as route reconvergence. Overall network topology is certainly a major factor for increasing uptime as are redundant circuits and paths. Also necessary are redundant power sources, additional CPUs and interfaces. Modular and self-healing routing software are also part of the solution.

Even with redundancy, some downtime or packet loss occurs when failover takes place from one circuit to the next, or between the primary hardware module and the backup system. Although the loss period should be minimal with redundant elements, failover can still be catastrophic for mission-critical, real-time applications. Reducing the recovery period is the focus of the latest HA routing efforts throughout the industry.

The key is to discover aspects of your network that perhaps you never knew.

## Truth in Testing

The truth about network performance can only be found with adequate testing. Validate the actual HA performance of a network or device by rigorously testing all facets of high availability routing. The only obstacle is that HA standards are still in draft format.

Rather than wait for IETF, some equipment manufacturers have developed proprietary or pre-standard solutions. When combined with prevailing standards, these implementations often yield satisfactory solutions for enterprises and service providers. The solutions must still be tested and quantified for users to fully understand the potential impact of network failure.

## Using the TeraRouting Tester

The three tests in this journal — VRRP/HSRP, BGP Graceful Restart and Graceful Restart for Other Protocols — are to be used with Spirent Communications’ TeraRouting Tester (TRT) product. The TRT application is utilized on Spirent’s SmartBits platform to emulate a large IP network running any major IPv4 or IPv6 routing protocol. After creating an emulated network, TRT generates traffic to and from all network devices and routes. A baseline for device capacity, throughput and latency can easily be established.

TRT simulates router problems or link failures as it generates traffic. The latency and packet loss that occur, as the router recovers, can be measured to help the user understand the “worst case” scenario for downtime and the data loss caused by failure. Much of the testing takes place from the data plane standpoint, the perspective that users have during network failure. While control plane events such as route updates are introduced in a controlled test environment, the relevant measurements are based on the loss or latency of data packets.

After a stable network is constructed in a lab, “negative” events can be introduced. Since TRT simulates router problems or link failures while generating traffic, the latency and packet loss that occurs during router recovery can be measured.

## Routing Protocols

All test methodologies in this journal use IPv4. However, these tests are equally applicable for IPv6 or for mixed IPv4/IPv6 networks. By changing addresses, the user can adapt the methodologies for an IPv6 environment. TRT fully supports IPv4 and IPv6, separately or simultaneously. Most routing protocols have extensions for graceful restart functionality (in official RFCs or specified in draft format RFCs).

## Explanation of the Test Methodologies

The three test methodologies are examples of testing solutions. Most use the BGP protocol, but the principles are equally applicable to other routing protocols. The tests require at least two SmartBits ports. More extensive testing can be conducted using the same basic methodology with additional ports.

In this document, Gigabit Ethernet ports are used for all tests. However, other media and speeds can be used. In most cases, the tests require one or two devices under test (DUTs) or routers. TRT modules generate the control plane infrastructure and the data plane traffic while TRT software measures and records results. The tests take advantage of TRT’s unique data plane clocking capabilities. Packets will be generated at a consistent rate of 1,000 per second.

Note: It is important that only data plane “signature” packets are counted, not the total TX/RX packet quantities since these include control plane information and may alter the totals. After a control plane event is generated, the resulting packet loss will be measured.

If 3,000 packets are lost, then one can infer the recovery latency was 3.000 seconds. If 18 packets were lost as a result of the negative event, the recovery time can be calculated to be an impressive speed of 0.018 seconds.

## Making A Difference

There is much more to know about testing high availability routers and reducing downtime. Spirent believes this journal is a good start because it provides you step-by-step testing procedures and basic educational material. If you would like to advance your knowledge in how to better control the effects of downtime, contact us for assistance.

Spirent Communications  
[www.spirentcom.com](http://www.spirentcom.com)

# Redundancy Protocols: VRRP and HSRP

## Associated RFCs

- RFC 3768 - Virtual Router Redundancy Protocol (VRRP) - April 2004
- RFC 2281 - Cisco Hot Standby Router Protocol (HSRP) - March 1998


## Overview

Router redundancy protocols facilitate the backup and recovery of an entire router or some of its major functional components. The Virtual Router Redundancy Protocol (VRRP) and the Hot Standby Router Protocol (HSRP) are similar in basic functionality.

VRRP is an industry standard protocol defined in RFC 3768. HSRP is generally applicable to Cisco environments, and it is specified in an informational RFC. In both cases, a primary router is selected and one or more backup device(s) is identified. A “hello” or “keep-alive” protocol operates between the two routers. If the secondary device detects a failure of the primary router (due to unanswered “hello” messages) it will assume the full identity of the primary router and take over its duties. Recovery times typically range from 10 seconds to two minutes.

## Objective

The redundancy protocol should operate between two routers with the test equipment generating, receiving and measuring traffic. Test the failover and recovery processes and latency for the two popular router redundancy protocols.



The diagram above indicates the basic test setup. (Similar test topologies will be used for the subsequent HA tests.) The SmartBits port on the right-hand side (Port-2) of the diagram will advertise BGP routes. The port on

the left (Port-1) will generate traffic streams at a periodic rate to all advertised networks. The primary router should forward this information to the SmartBits port (Port-2) on the right. After a router failure occurs (i.e., by turning the primary router off) the backup router should assume the traffic forwarding functions. The SmartBits port on the right (Port-2) will measure the latency and packet loss associated with this switch over.

**Note:** The Ethernet segments are depicted in the diagram as old-fashioned LAN segments. In reality, a gigabit hub or switch is more likely to be used. *It is critical that a hub be used instead of a switch.* Otherwise the test results will be skewed to include the additional latency incurred while the switch updates its forwarding table to reflect the changeover from the primary router to the backup router. It is also highly recommended that the hub be dedicated to this particular test and not support any other extraneous traffic.

### Step-by-Step Testing

1. Configure the routers for VRRP or HSRP. Identify the primary and secondary routers, and also assign the appropriate virtual IP addresses. Make sure the configuration includes a peer address and Autonomous System Number (ASN) for the BGP router that will be emulated on TRT Port-2.
2. Configure the TRT SmartBits Port-1 to “Xmit Only.” Select “Xmit and Advertise” (or “Advertise Only”) for Port-2. Enable BGP routing on Port-2.

Show Columns

Port Info   
  Hub,Slot,Port   
  Ethernet   
  POS   
  POS-PPP  
 Alias   
 IPv4   
 IPv6   
 ATM

Drag a column header here to group by that column.

Port	Link	Alias	Cluster	VPN ID	Model	Participation	Reservation ▲	Capture	VLAN	Static	BGP
bills_smb 2-1		Port-1			LAN-3325A	Xmit Only	Reserved	Off	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bills_smb 2-2		Port-2			LAN-3325A	Xmit and Advertise	Reserved	Off	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
bills_smb 2-3					LAN-3325A	Not Used	Reserved	Off	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
▶ bills_smb 2-4					LAN-3325A	Not Used	Reserved	Off	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

- Set the correct IP addresses for the SmartBits ports. Note that the gateway addresses should be the Virtual IP addresses for the VRRP or HSRP groups.

Show Columns

Port Info     Hub,Slot,Port     Ethernet     POS     POS-PPP  
 Alias     IPv4     IPv6     ATM

Drag a column header here to group by that column.

Port	IPv4 Address	IPv4 Gateway	IPv4 Prefix Length	MAC Address
bills_smb 2-1	10.0.0.2	10.0.0.1	24	00-00-02-00-00-01
bills_smb 2-2	20.0.0.2	20.0.0.1	24	00-00-02-00-00-02
bills_smb 2-3	192.91.1.2	192.91.1.1	24	00-00-02-00-00-03
▶ bills_smb 2-4	192.92.1.2	192.92.1.1	24	00-00-02-00-00-04

- Enable a single BGP peer on Port-2. Note that the address and ASN must match those that are already configured in the DUTs.

Untitled - TeraRouting Tester

File Edit View Tools Actions Help

Apply Result Collection... Wizards...

Test

- Card Setup
- Sub-interface
- Static Route
- BGP
  - Sessions
  - Routes
  - Options
- OSPF
- OSPFv3
- RIP
- IS-IS

Apply Default Addresses

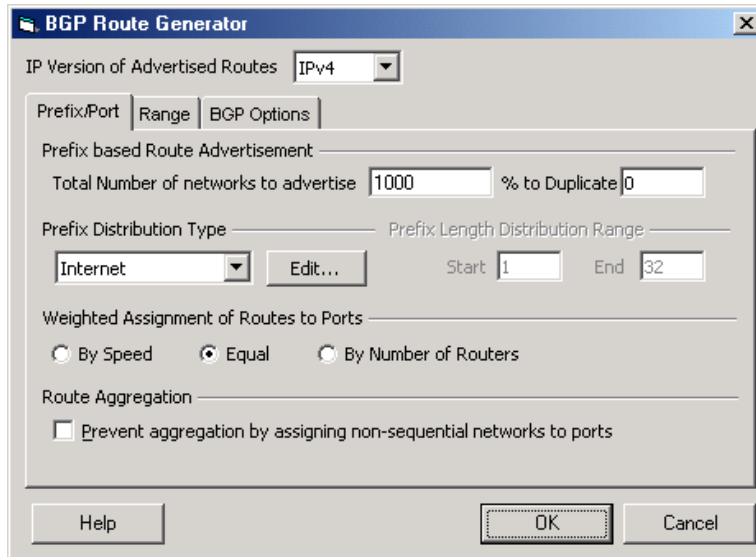
Show columns

IPv6     VLAN     ATM     Intervals    BGP Router ID: Same as Tester IPv4

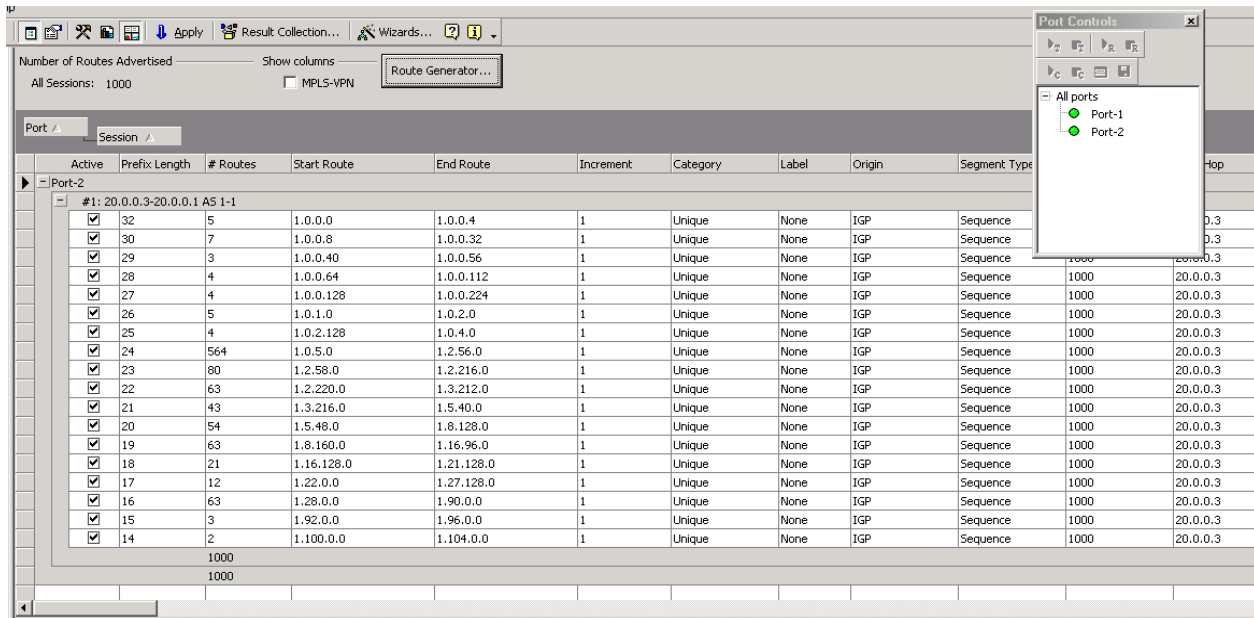
Drag a column header here to group by that column.

Port	State	Tester AS #	SUT AS#	I/E-BGP	Session IPv	Route IPv	Tester IPv4 Address	IPv4 Prefix Length
▶ Port-2	None	1	1	I-BGP	IPv4	IPv4	20.0.0.3	24

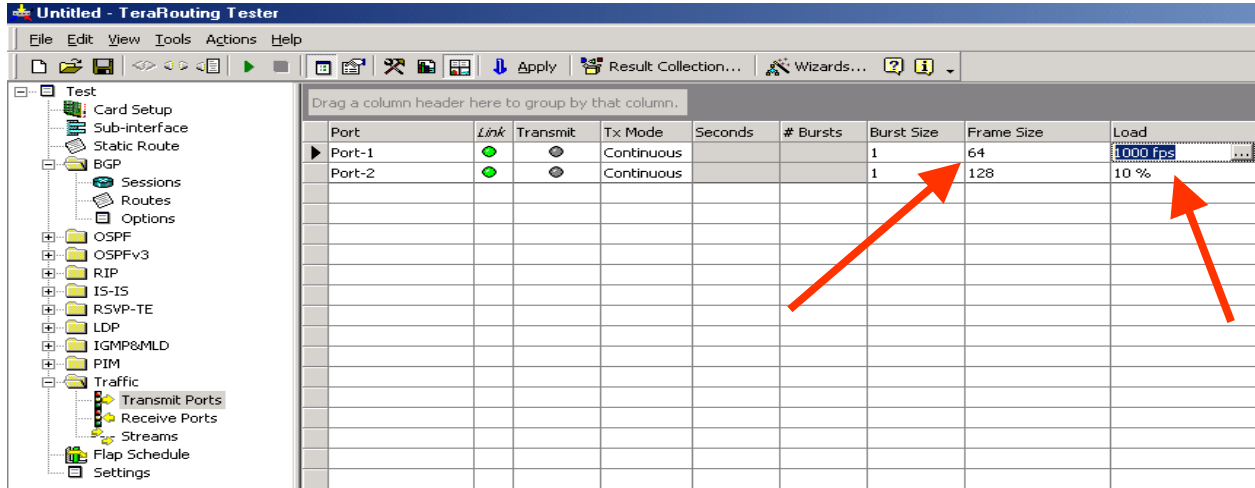
- Use the BGP Route generator to create some quantity of IPv4 routes that will be advertised to the DUTs. The quantity and distribution types are probably not significant, but these can be varied in subsequent iterations of the test.



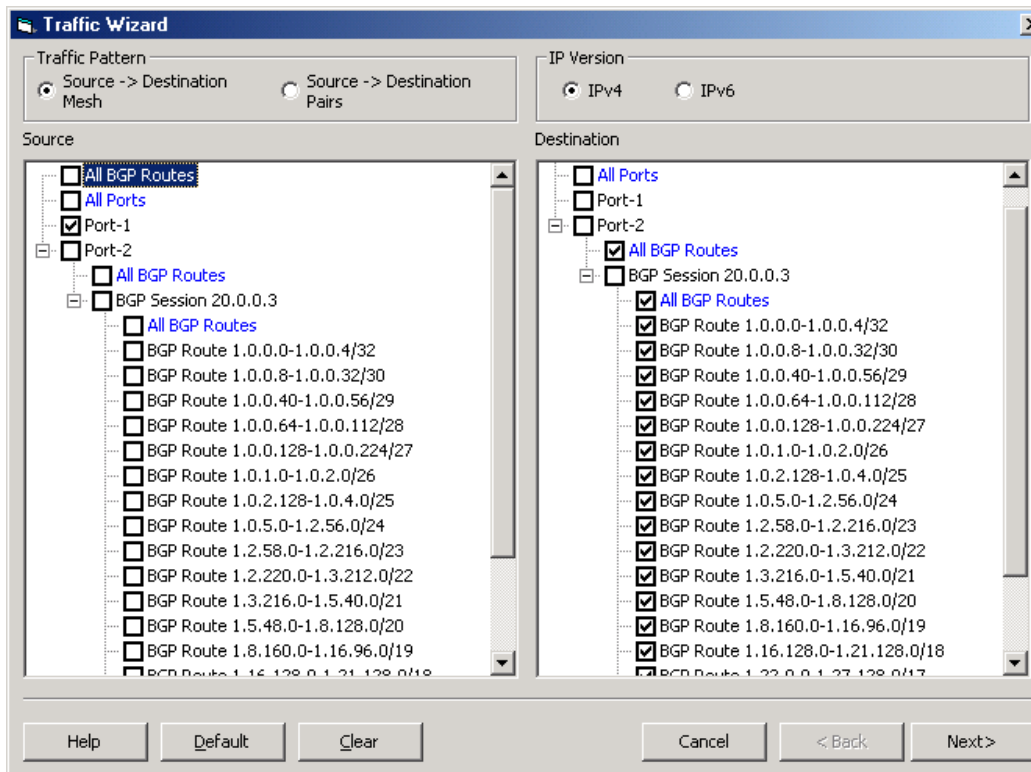
- The resulting routing table on TRT for 1,000 routes with an Internet distribution should look like this:



- Next, go to the “Transmit Ports” configuration screen. Set the frame size for Port-1 to 64 bytes and set the load value to 1,000 frames per second.



- Use the “Streams Wizard” to generate traffic from Port-1 to “All BGP Routes.”



9. The resulting traffic streams table should look like this:

Active	Name					Rx Port(s)	Fill Pattern
<input checked="" type="checkbox"/>	Stream1	Eth	IPv4 DA	1.0.0.0-1.0.0.4/32,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream2	Eth	IPv4 DA	1.0.0.10-1.0.0.34/30,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream3	Eth	IPv4 DA	1.0.0.46-1.0.0.62/29,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream4	Eth	IPv4 DA	1.0.0.78-1.0.0.126/28,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream5	Eth	IPv4 DA	1.0.0.158-1.0.0.254/27,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream6	Eth	IPv4 DA	1.0.1.62-1.0.2.62/26,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream7	Eth	IPv4 DA	1.0.2.254-1.0.4.126/25,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream8	Eth	IPv4 DA	1.0.5.3-1.2.56.3/24,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream9	Eth	IPv4 DA	1.2.58.3-1.2.216.3/23,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream10	Eth	IPv4 DA	1.2.220.3-1.3.212.3/22,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream11	Eth	IPv4 DA	1.3.216.3-1.5.40.3/21,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream12	Eth	IPv4 DA	1.5.48.3-1.8.128.3/20,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream13	Eth	IPv4 DA	1.8.160.3-1.16.96.3/19,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream14	Eth	IPv4 DA	1.16.128.3-1.21.128.3/18,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream15	Eth	IPv4 DA	1.22.0.3-1.27.128.3/17,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream16	Eth	IPv4 DA	1.28.0.3-1.90.0.3/16,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream17	Eth	IPv4 DA	1.92.0.3-1.96.0.3/15,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
<input checked="" type="checkbox"/>	Stream18	Eth	IPv4 DA	1.100.0.3-1.104.0.3/14,SA	192.89.1.2/32,GW	192.89.1.1	All Zeros
Total:18							

10. Run the test. The remaining default settings are acceptable.

Test Configuration

Tests | Test Options | Capture

Test Type: Continuous

Duration:  Seconds: 10  Bursts: 1000


Load Options: Load Units: Percent (%)

Frame Size:  Use Transmit Ports/Streams Setting  Fixed: 128  Random: Min: 128 Max: 256

Load:  Use Transmit Ports/Streams Setting  Fixed: 10  Random: Min: 10 Max: 50

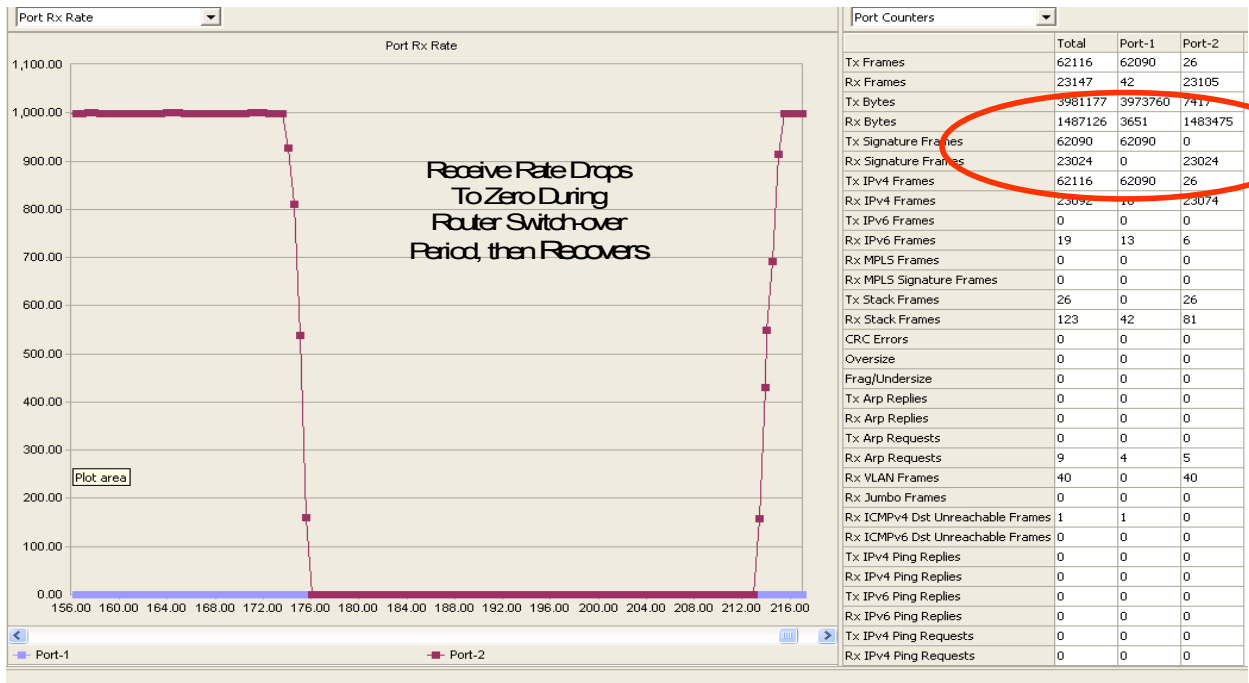
Run OK Cancel

11. Validate all of the test parameters. The traffic, as indicated below, should be received on Port-2 at a consistent rate of 1,000 frames per second.



12. If necessary, clear all of the counters by right-clicking on the “Port Counters” chart.
13. Induce a failure of the primary router (powering the device off is an easy way to accomplish this).
14. Watch the graphs to see the forwarding rate go to zero, and then recover.

15. Calculate the recover or switch-over time interval. The graph below suggests the recover time was approximately 40 seconds. A more precise calculation can be derived from the counters on the right-hand side. 62,090 signature frames were transmitted from Port-1. 23,024 frames were received on Port-2. The difference between these quantities is 39,066 frames. Since Port-1 is transmitting at a consistent rate of 1,000 frames per second, this indicates that the “downtime” was 39.066 seconds, which is consistent with the graphical estimate.



16. This test can be repeated using different protocols, traffic patterns, routing tables and port quantities.

# BGP Graceful Restart

## Associated RFCs

- draft-ietf-idr-restart-10 - Graceful Restart Mechanism for BGP - June 2004
- RFC 3623 - Graceful OSPF Restart - November 2003
- RFC 3478 - Graceful Restart Mechanism for Label Distribution Protocol - February 2003
- RFC 3847 - Restart Signaling for Intermediate System to Intermediate System (IS-IS) - July 2004

## Overview

Graceful restart mechanisms were first developed for the BGP protocol due to constant changes within the Internet. The latest graceful restart draft proposal is Draft 10 from June 2004 and is posted on the IETF's Web site. The document name is draft-ietf-idr-restart-10.

## Objective

This test validates the device's BGP graceful restart functionality and measures the recovery times. In this test, TRT participates directly in the graceful restart protocol.

Graceful restart is an optional extension to BGP that allows routers to "dampen" the effect of "flapping" peers. If a BGP router fails, then recovers, it will retransmit its entire routing table, causing all of its peers to update and recalculate their routing tables. If a router is "flapping," the resulting control plane overhead will be quite significant as all of the adjacent routers constantly recalculate their respective forwarding information bases. If graceful restart is used, these calculations will not take place until the configured restart timer expires (typically 90 seconds). This delay reduces the impact of "flapping" on the adjacent routers, and it also allows for a hold period during which stability can hopefully be restored. If the router recovers prior to the expiration of the hold timer, the recovery will be quite rapid since the adjacent router will not need to recalculate its routing database.


Since a "flapping" router might only transmit a portion of its routing table before it changes its state again, an end-of-RIB (routing information base) marker is used to indicate the end of the database. A router that supports graceful restart will not update its own routing table until this end-of-RIB marker is received, thereby indicating that a complete database has been successfully transmitted. This process is especially useful when the routers are attached to the Internet and the resulting databases contain well over 100,000 entries.

TRT supports the graceful restart extension of BGP, so it is possible to natively test this functionality. The test setup requires a single router and two SmartBits/TRT ports. Port-1 will be used for traffic generation. Port-2 will be used to establish a BGP session with the DUT; the graceful restart extension will also be invoked.

After the BGP sessions are established and traffic is being forwarded at a fixed rate (1,000 packets per second), a control plane event can be introduced. One such event could consist of breaking the underlying TCP session between the DUT and the emulated BGP router. In response to this failure, the DUT will cease forwarding traffic to the next hop router. However, if that router is restored prior to the expiration of the hold timer, the resulting recovery should be quite rapid.

For the example and instructions below, some exaggerated timers and counters will be used. These can be modified based on the required parameters of each specific DUT.

## Setup



## Step-by-Step Testing

1. Configure the DUT. BGP should be enabled on Port-2 with a peer defined for the SmartBits port. Enable Graceful Restart for that connection.
2. Configure the TRT SmartBits Port-1 to "Xmit Only." Select "Xmit and Advertise" (or "Advertise Only") for Port-2. Enable BGP routing on Port-2.

Show Columns

Port Info     Hub,Slot,Port     Ethernet     POS     POS-PPP  
 Alias     IPv4     IPv6     ATM

Drag a column header here to group by that column.

Port	Link	Alias	Cluster	VPN ID	Model	Reconfiguration	Reservation	Capture	VLAN	Stack	BGP	OSPF	OSPFv3
bills_smb 2-1		Port-1			LAN-3325A	Xmit Only	Reserved	Off	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bills_smb 2-2		Port-2			LAN-3325A	Xmit and Advertise	Reserved	Off	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
bills_smb 2-3					LAN-3325A	Not Used	Reserved	Off	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
▶ bills_smb 2-4					LAN-3325A	Not Used	Reserved	Off	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

- Set the correct IP addresses for the SmartBits ports.

Table showing SmartBits ports configuration:


Port	IPv4 Address	IPv4 Gateway	IPv4 Prefix Length	MAC Address	SUT MAC
bills_smb 2-1	192.168.100.2	192.168.100.1	24	00-00-02-00-00-01	00-00-00-
bills_smb 2-2	192.168.200.2	192.168.200.1	24	00-00-02-00-00-02	00-00-00-
bills_smb 2-3	192.168.200.3	192.168.200.1	24	00-00-02-00-00-03	00-00-00-
bills_smb 2-4	192.92.1.2	192.92.1.1	24	00-00-02-00-00-04	00-00-00-

- Enable a single BGP peer on Port-2. Note that the address and ASN must match those that are already configured in the DUT.

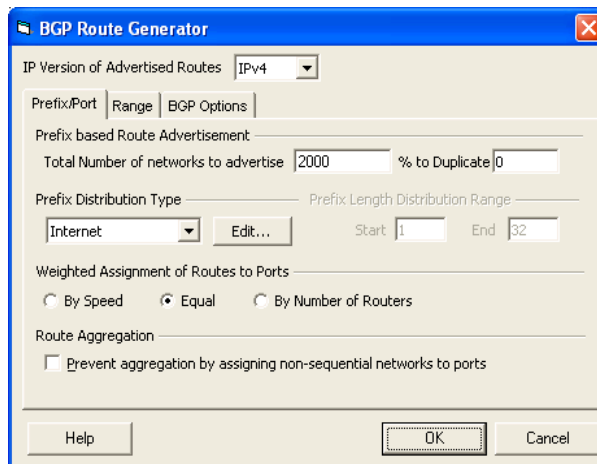
Table showing BGP configuration for Port-2:

Port	State	Tester AS #	SUT AS#	I/E-BGP	Session IPv	Route IPv	Tester IPv4 Address	IPv4 Prefix Length
Port-2	None	1	20	E-BGP	IPv4	IPv4	192.168.200.3	24

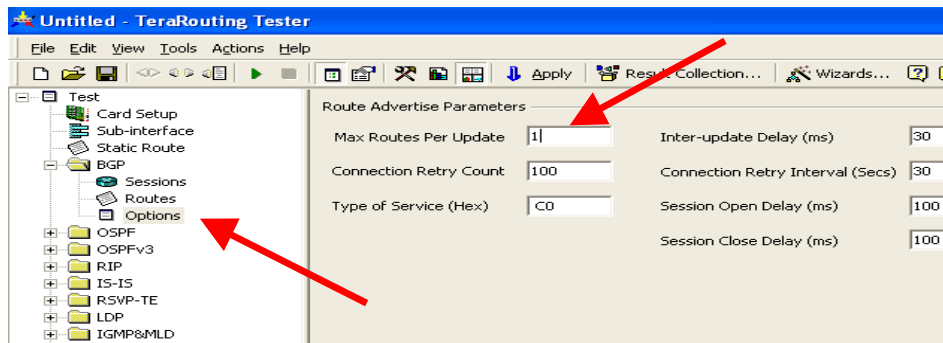
- On the same configuration screen, scroll to the right-hand side and check the box to “Enable Restart.” Select the appropriate value for the restart timer (should match the DUT configuration).



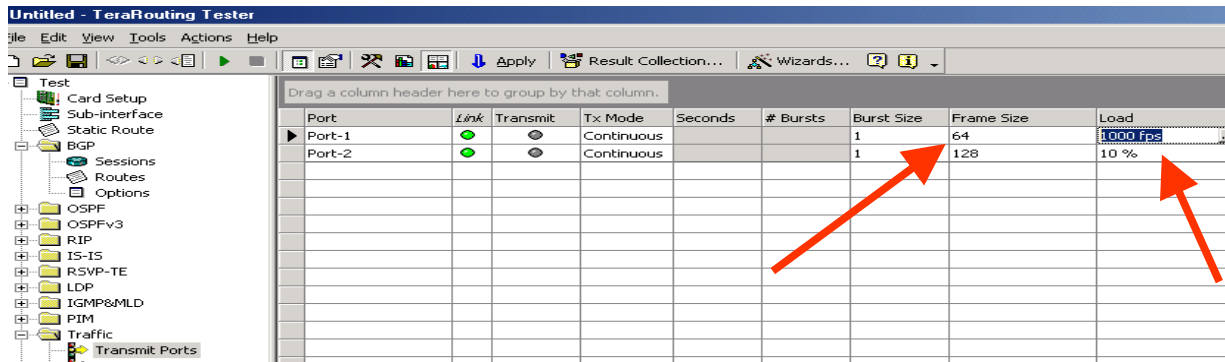
- Use the BGP Route Generator to advertise 2,000 routes. Although the prefix distribution method probably is not significant for this test, the example below uses an Internet distribution.



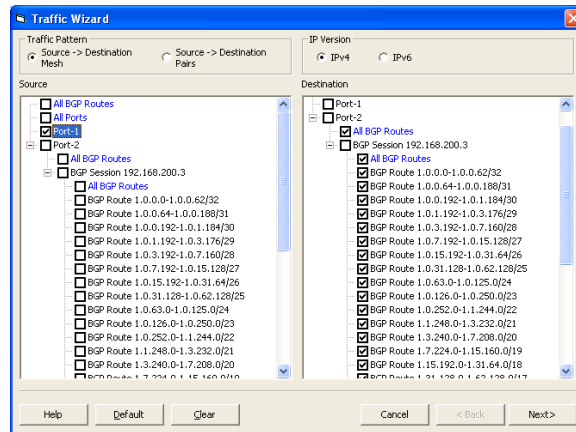
7. Select the BGP Options configuration screen from the navigation window. Configure the “Max Routes Per Update” for a value of 1 – this is a painfully slow rate since BGP will typically advertise hundreds of routes in a single update. However, this will significantly prolong the route advertisement times, allowing easy verification of the graceful restart functionality. If it is necessary to further reduce the rate of BGP advertisements (thereby increasing the associated convergence times), the “Inter-Update Delay” can also be increased.



8. Next, go to the “Transmit Ports” configuration screen. Set the frame size for Port-1 to 64 bytes and set the load value to 1,000 frames per second.



9. Use the “Streams Wizard” to generate traffic from Port-1 to “All BGP Routes.”



10. The resulting 19 stream definitions should look like this:

- Test
- Card Setup
- Sub-interface
- Static Route
- BGP
  - Sessions
  - Routes
  - Options
- OSPF
  - OSPFv3
- RIP
- IS-IS
- RSVP-TE
- LDP
- IGMP&MLD
- PIM
- Traffic
  - Transmit Ports
  - Receive Ports
  - Streams
  - Flap Schedule
  - Settings

Traffic Configuration

Streams

Traffic Wizard...

Traffic Groups...

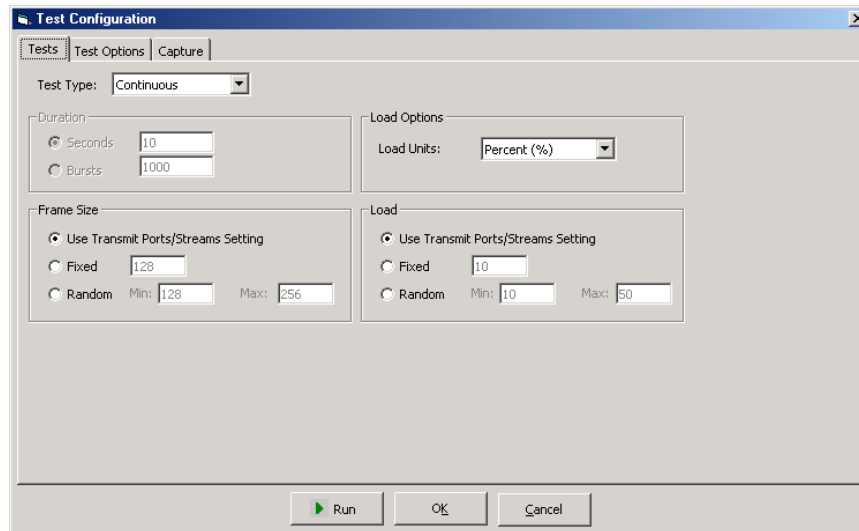
Edit...

Total Streams: 19

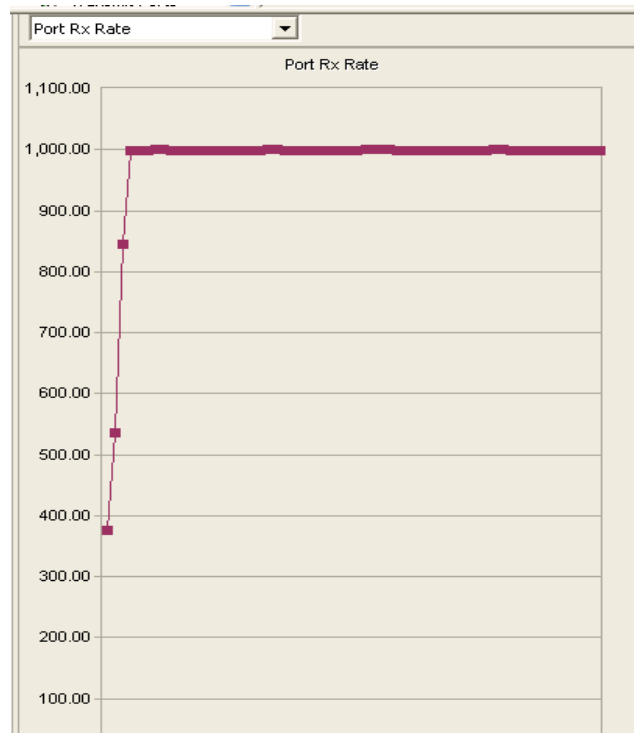
Active Streams: 19

Port	Active	Name		Rx Port(s)	Fill Pattern	
bills_smb 2-1	<input checked="" type="checkbox"/>	Stream1	Eth	IPv4 DA 1.0.0.0-1.0.0.10/32,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream2	Eth	IPv4 DA 1.0.0.14-1.0.0.70/30,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream3	Eth	IPv4 DA 1.0.0.78-1.0.0.126/29,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream4	Eth	IPv4 DA 1.0.0.142-1.0.0.254/28,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream5	Eth	IPv4 DA 1.0.1.30-1.0.1.254/27,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream6	Eth	IPv4 DA 1.0.2.62-1.0.4.126/26,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream7	Eth	IPv4 DA 1.0.4.254-1.0.8.254/25,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream8	Eth	IPv4 DA 1.0.9.3-1.4.102.3/24,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream9	Eth	IPv4 DA 1.5.172.3-1.7.164.3/22,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream10	Eth	IPv4 DA 1.4.104.3-1.5.168.3/23,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream11	Eth	IPv4 DA 1.7.168.3-1.10.80.3/21,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream12	Eth	IPv4 DA 1.10.96.3-1.17.16.3/20,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream13	Eth	IPv4 DA 1.17.32.3-1.32.192.3/19,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream14	Eth	IPv4 DA 1.33.0.3-1.43.128.3/18,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream15	Eth	IPv4 DA 1.44.0.3-1.55.128.3/17,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream16	Eth	IPv4 DA 1.56.0.3-1.182.0.3/16,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream17	Eth	IPv4 DA 1.184.0.3-1.196.0.3/15,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream18	Eth	IPv4 DA 1.200.0.3-1.212.0.3/14,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
	<input checked="" type="checkbox"/>	Stream19	Eth	IPv4 DA 1.216.0.3/13,SA 192.89.1.2/32,GW 192.89.1.1	Any	All Zeros
Total:19						

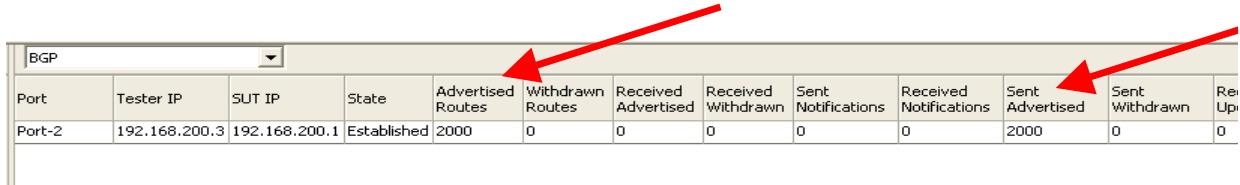
11. Run the test. The remaining default settings are acceptable.



12. Validate all of the test parameters. The traffic, as indicated below, should be received on Port-2 at a consistent rate of 1,000 frames per second. Note that there will be a significant delay prior to the start of traffic (60 seconds or longer) while TRT advertises all 2,000 routes at a rate of one advertisement every 30 milliseconds (or 33 routes per second).



13. Validate the quantity of routes advertised by using the BGP Stats table. Also validate the quantity of routes the DUT received. In both cases, the total should be 2,000.

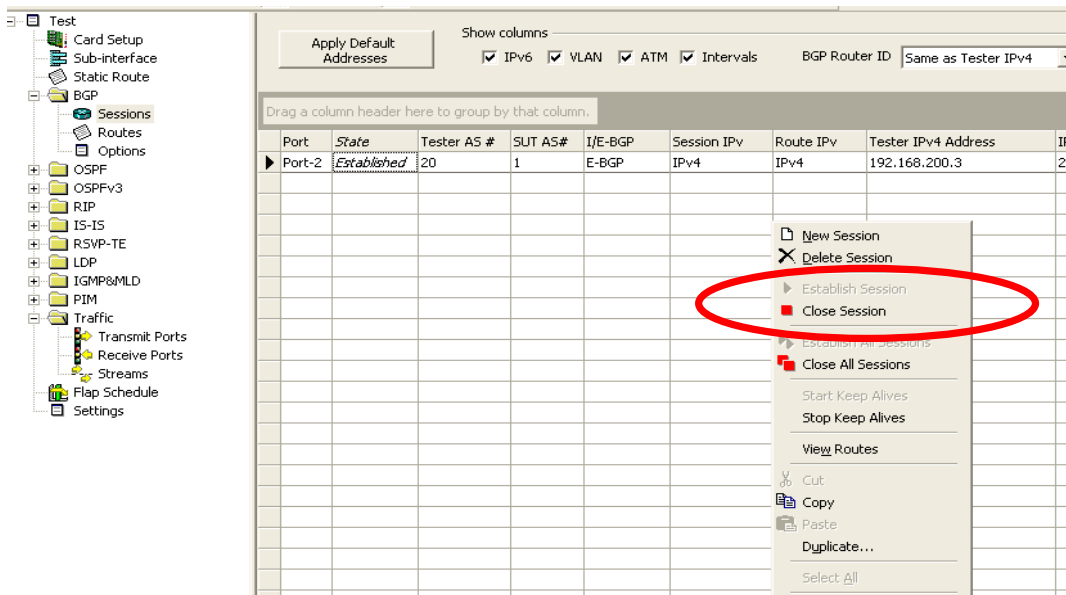


Port	Tester IP	SUT IP	State	Advertised Routes	Withdrawn Routes	Received Advertised	Received Withdrawn	Sent Notifications	Received Notifications	Sent Advertisements	Sent Withdrawn	Re Up
Port-2	192.168.200.3	192.168.200.1	Established	2000	0	0	0	0	0	2000	0	0

14. If necessary, clear all of the counters by right-clicking on the “BGP Counters” chart.

15. Introduce a negative event. There are two ways to accomplish this objective (select one of the following bulleted options below):

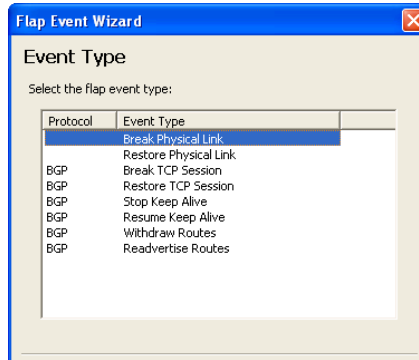
- Return to the “BGP Sessions” configuration menu and right-click on the established session. This will invoke a pop-up menu with several new options. Select the option to “Close Session” – this will close the existing BGP session.



The screenshot shows the BGP Sessions configuration interface. On the left is a tree view with categories like BGP, OSPF, RIP, etc. The main area displays a table of sessions. A context menu is open over the first session, with the 'Close Session' option highlighted by a red oval.

Port	State	Tester AS #	SUT AS#	I/E-BGP	Session IPv	Route IPv	Tester IPv4 Address	IP
Port-2	Established	20	1	E-BGP	IPv4	IPv4	192.168.200.3	24

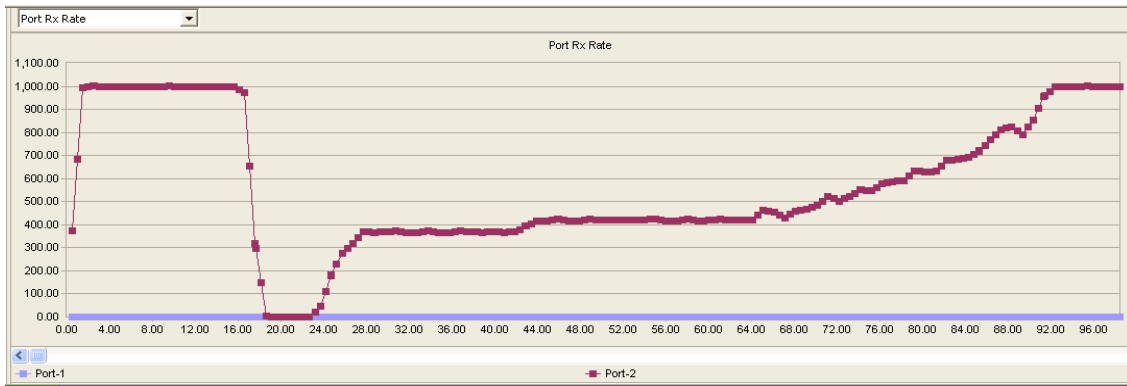
- The other alternative is to use the “Flap Scheduler.” This cannot be used interactively, and can only affect “batch mode” tests. However, it will provide many more selections for negative events. The BGP events are indicated in the screenshot below. It is advisable to test all of these permutations.




16. Shortly after creating the “negative event,” restore the configuration to its previous working state by counteracting the event (re-establish the BGP session, for example). This must be accomplished before the expiration of the restart timer (default is 90 seconds). Note the resulting increase in the quantity of advertised routes, since the restored peer must send 2,000 new advertisements. This process will take another 60 seconds or so.

Port	Tester IP	SUT IP	State	Advertised Routes	Withdrawn Routes	Received Advertised	Received Withdrawn	Sent Notifications	Received Notifications	Sent Advertised	Sent Withdrawn
Port-2	192.168.200.3	192.168.200.1	Established	4000	0	0	0	0	0	4000	0

17. A router that does not have graceful restart capabilities will need to wait until all of the advertisements are complete in order to start forwarding traffic. However, a router that uses graceful restart should be able to resume forwarding data as soon as the peer (or physical link) is restored. The graceful restart recovery time is expected to be a few seconds instead of a minute. See comparison below for a given (and nameless) router.



**Without Graceful Restart  
Full Recovery Takes 76 Seconds**



**With Graceful Restart,  
Full Recovery Takes 18  
Seconds**

18. These tests can be repeated while varying traffic loads, route quantities, port quantities and other control plane parameters. Also try selecting several other types of negative events.

# Graceful Restart for Other Protocols

## Associated RFCs

- draft-ietf-idr-restart-10 - Graceful Restart Mechanism for BGP - June 2004
- RFC 3623 - Graceful OSPF Restart - November 2003
- RFC 3478 - Graceful Restart Mechanism for Label Distribution Protocol - February 2003
- RFC 3847 - Restart Signaling for Intermediate System to Intermediate System (IS-IS) - July 2004

## Overview

Although graceful restart is most applicable for the Internet environment and the BGP protocol, the IETF is working on developing graceful restart mechanisms for the other protocols. Graceful restart extensions have been proposed for OSPF, IS-IS and LDP. Since principles of graceful restart are very similar for all routing protocols, the test methodologies will be comparable.

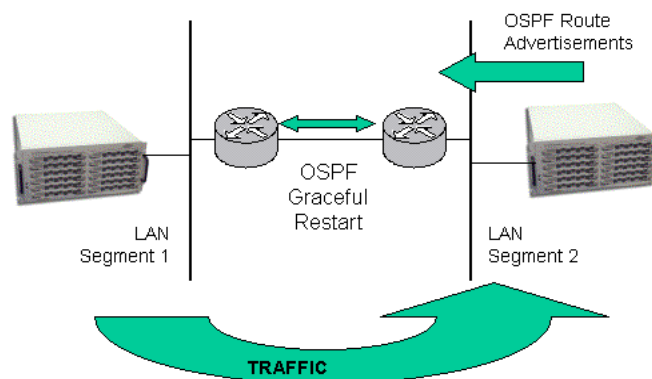
## Objective

This test validates the device's graceful restart functionality for all protocols and measures the recovery times.

TRT currently only supports graceful restart for BGP (the most popular HA routing protocol). Although TRT cannot directly participate in the graceful restart functionality for LDP, OSPF and IS-IS it can still effectively test and validate the recovery mechanisms of these protocols. These tests will simply be conducted from the data plane's perspective.

For this example, OSPF will be used for this test. However, this methodology can be applied to LDP or IS-IS. In fact, BGP graceful restart can also be tested in this manner.

## Setup

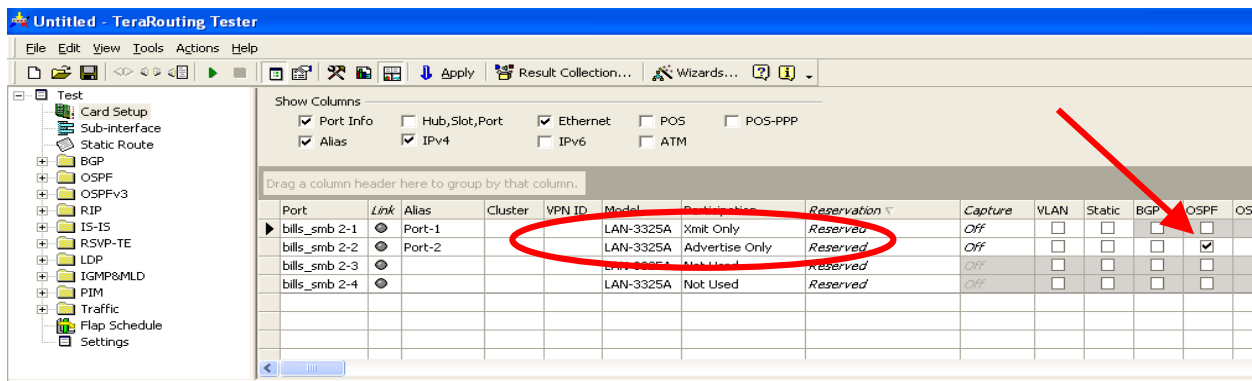


The diagram indicates the test setup for an OSPF Graceful Restart test. Two routers are required, along with two (or more) TRT/SmartBits ports. OSPF adjacencies should be established between the two routers and Port-2 (on the right-hand side of the diagram) of the SmartBits tester. It is suggested that all of these ports are initially placed into the same OSPF area (in the example, we will use the backbone area 0.0.0.0). However, inter-area testing can be conducted.

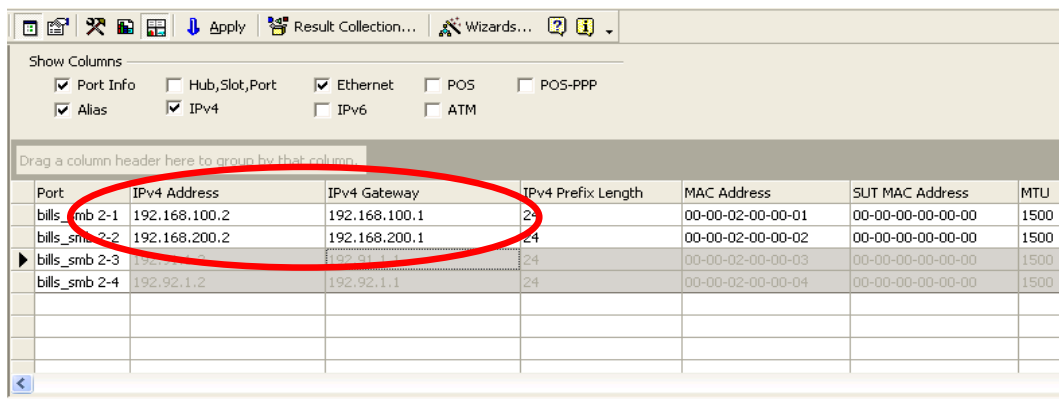
The TRT application will generate OSPF Link State Advertisements (LSAs) for various networks, links and devices. The routers will then exchange this information. Port-1 of the SmartBits tester will then generate traffic to all of these advertised networks. After validating the successful receipt of traffic, an error condition will be introduced between the two routers causing the traffic to stop. The error will then be corrected and the recovery time measured. The graceful restart function should significantly reduce the recovery time period.

### Step-by-Step Testing for OSPF Graceful Restart

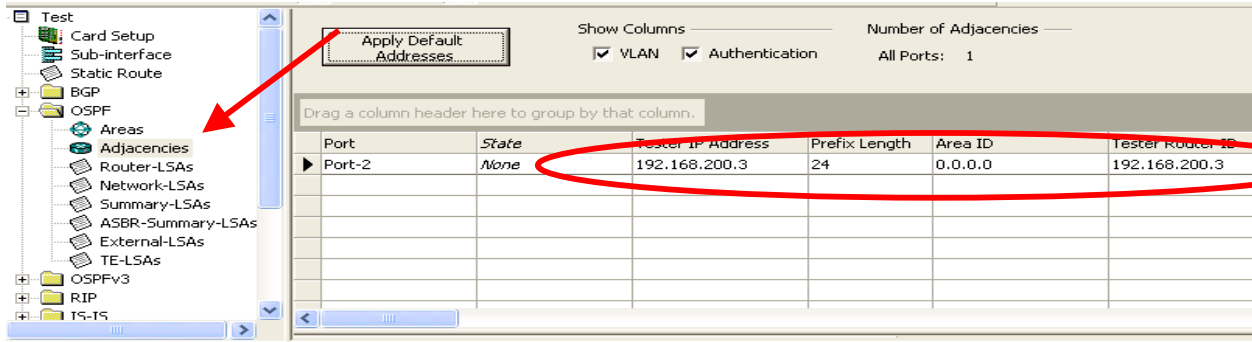
1. Configure the DUTs. OSPF should be enabled on both devices using area 0.0.0.0. Graceful restart should be enabled on the ports connecting the two routers (not on the ports connecting to the SmartBits testers).
2. Configure the TRT SmartBits Port-1 for “Xmit Only.” Select “Xmit and Advertise” (or “Advertise Only”) for Port-2. Enable OSPF routing on Port-2.



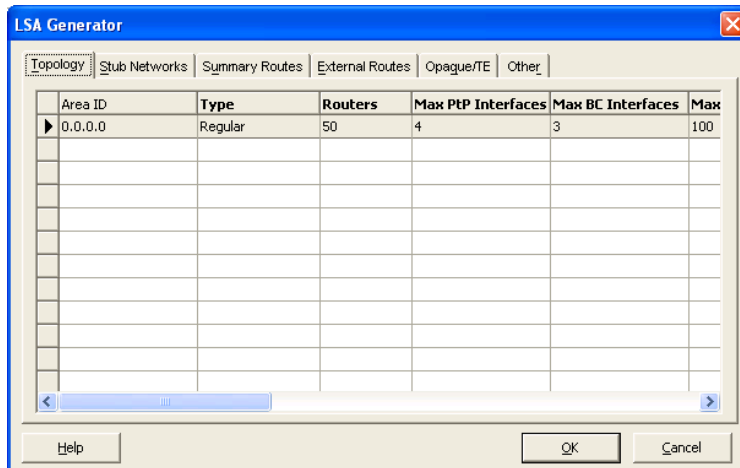
3. Set the correct IP addresses on the SmartBits ports.



- Go to the OSPF Adjacencies configuration menu. Set the tester's IP address, Router ID and Area ID.



- Use the LSA Generator to produce a reasonable mix of Link State Advertisements. For most purposes, the default values should be sufficient. Although larger networks will yield more dramatic results for graceful restart measurements.



6. The resulting topology should look like this:

The screenshot displays the TeraRouting Tester interface. At the top, there is a table for 'View LSAs Advertised In' for Area 0.0.0.0. Below the table is a network diagram showing a central router (192.168.200.3) connected to multiple other routers. A summary window titled 'LSA Generator Summary' is open on the right, showing statistics for the area.

Advertising Router	Options	Router Type	# Links	Link Type	Link ID	Link Data	Metric
192.168.200.3	02	EB	7	PTP	0.0.0.1	0.0.0.1	1
				PTP	0.0.0.3	0.0.0.2	1
				PTP	0.0.0.5	0.0.0.3	1
				PTP	0.0.0.7	0.0.0.4	1
				Transit Network	1.0.175.1	1.0.175.1	1
				Transit Network	1.0.176.1	1.0.176.1	1
				Transit Network	1.0.177.1	1.0.177.1	1
0.0.0.1	02		4	PTP	0.0.0.17	0.0.0.1	1
				PTP	0.0.0.49	0.0.0.2	1
				Transit Network	1.0.178.1	1.0.178.1	1
				PTP	192.168.200.3	0.0.0.3	1

**LSA Generator Summary**

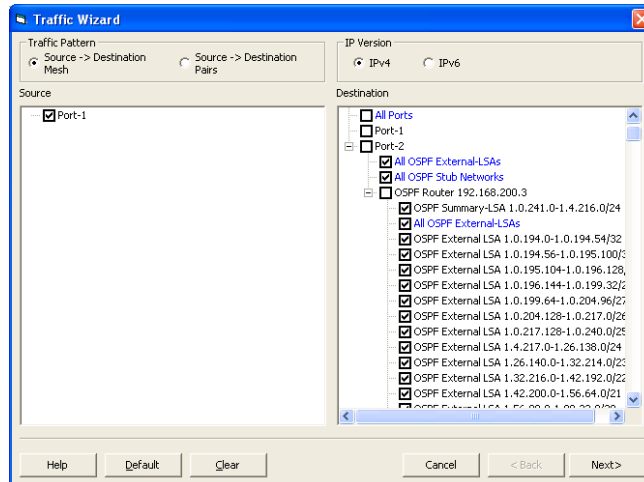
Area = 0.0.0.0  
 Router-LSAs: 50  
 Stub Networks: 300  
 Network-LSAs: 19  
 Summary-LSAs: 1000  
 Duplicated Summary-LSAs: 0  
 AS-External/NSSA LSAs: 10000  
 Duplicated AS-External/NSSA LSAs: 0  
 TE LSA with Router Address TLV: 0  
 TE LSA with Link TLV: 0

7. Go to the “Transmit Ports” configuration screen. Set the frame size for Port-1 to 64 bytes, and set the load value to 1,000 frames per second.

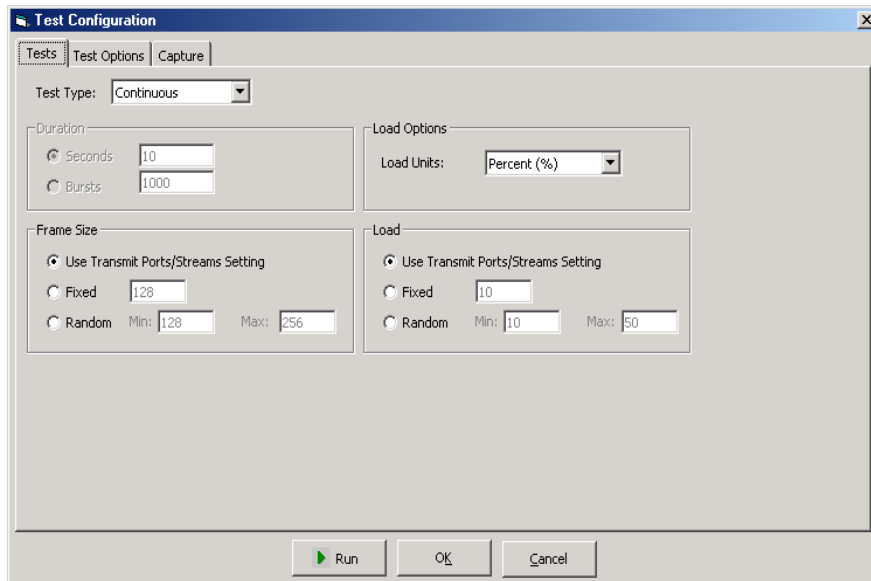
The screenshot shows the 'Transmit Ports' configuration window in TeraRouting Tester. The window title is 'Untitled - TeraRouting Tester'. The 'Test' pane on the left shows a tree view with 'Card Setup', 'Sub-interface', 'Static Route', 'BGP', and 'OSPF'. The main configuration table is as follows:

Port	Link	Transmit	Tx Mode	Seconds	# Bursts	Burst Size	Frame Size /	Load
Port-1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Continuous			1	64	1000 fps


- Use the “Streams Wizard” to generate traffic from Port-1 to “All OSPF External-LSAs” and “All OSPF Stub Networks.”



- If you selected the default settings for the OSPF LSA Creator, a total of 323 streams should be generated. This quantity will vary based on the total numbers of emulated routers and networks that you selected.
- Run the test. The remaining default settings are acceptable.




11. Validate all test parameters. The traffic, as indicated below, should be received on Port-2 at a consistent rate of 1,000 frames per second.



12. Clear the counters if necessary. To do this, simply right-click on the “Port Counters” chart.
13. Introduce a control plane failure between the two routers. This can be accomplished several different ways. The cable between both routers can be disconnected, OSPF can be disabled on one of the devices, an interface can be disabled, etc.

- Repair the failure and measure the recovery time. This step should be repeated with and without graceful restart enabled so that the difference between the two recovery times can be measured. Graphically, the diagram below suggests the recovery time is approximately 33 seconds. The calculated packet loss confirms this visual estimate.



- This test can be repeated using different protocols, port counts, traffic patterns and routing tables. The failure mechanism should also be varied on subsequent iterations to determine the impact on the recovery latency.

# Glossary

This glossary provides definitions of the terminology, abbreviations, and acronyms used in this Spirent Communications test methodology journal as well as terms based on telecom technology and standards.

## ***Fast Re-Route.***

An MPLS protection mechanism. This requires the creation of a standby backup LSP using a different path than the primary circuit. If the primary LSP fails, the recovery process involves a rapid failover to the secondary link.

## ***Five Nines.***

A communications industry buzz word (first popularized by the telephony industry) meaning the objective for communications equipment is 99.999 percent uptime.

## ***High Availability Routing (HA Routing).***

An industry term referring to the ability for routers and networks to provide nearly non-stop services. Many facets (redundancy, failover and recovery mechanisms, etc.) help support HA routing.

## ***Hitless Upgrades.***

A means for upgrading a network device (hardware and/or software) without impacting network traffic. In other words, a router can continue to forward packets while a new interface, a new revision of software or a new configuration file is activated.

## ***Hot Standby Router Protocol (HSRP).***

A redundancy and failover protocol that occurs between routers. This is primarily used by Cisco and is defined in an informational RFC. This protocol is similar to VRRP.

## ***Non-Stop Forwarding.***

An industry term nearly synonymous with High Availability Routing. The subtle difference is that Non-Stop Forwarding implies that a routing paradigm is implemented which tolerates absolutely no downtime.

## ***Virtual Router Redundancy Protocol (VRRP).***

A redundancy and failover protocol that occurs between routers. VRRP is a “standard” method defined by the IETF in RFC 3768 and similar to Cisco's Hot Standby Router Protocol (HSRP).

# Spirent Communications Test Methodologies Information

Spirent Communications at <[www.spirentcom.com](http://www.spirentcom.com)> has a library of test methodology journals that offer you insight into testing a variety of technologies.

This issue of the Spirent Communications Step-by-Step Test Methodologies journal uses the TeraRouting Tester (TRT) on the SmartBits platform.

The following platform-specific journals can be found on Spirent's Customer Service Center (CSC) at <<http://support.spirentcom.com>>. Registration is required.

- TRT Edition: High Availability Routing
- TRT Edition: IGMP/MLD Functional Testing
- TRT Edition: Protocol Independent Multicast (PIM)

While this journal is platform-specific for TRT, Spirent also provides a downloadable series of generic and informative test methodology journals at <[www.spirentcom.com/tmj](http://www.spirentcom.com/tmj)>:

- Router Performance Edition
- Volume II - Wireless, Edge Router, Metro Optical, VoIP, SS7, QoS and IPv6
- IPv6 Edition
- Multicast Edition
- PPPoX Edition
- Optical Edition
- Edge Router Edition
- MPLS Edition
- IPSec Edition
- Layer 4-7 Edition

Sales and technical support phone numbers are listed on the front inside cover of this journal in case you would like to contact us.



*Analyze* | *Assure* | *Accelerate™*

