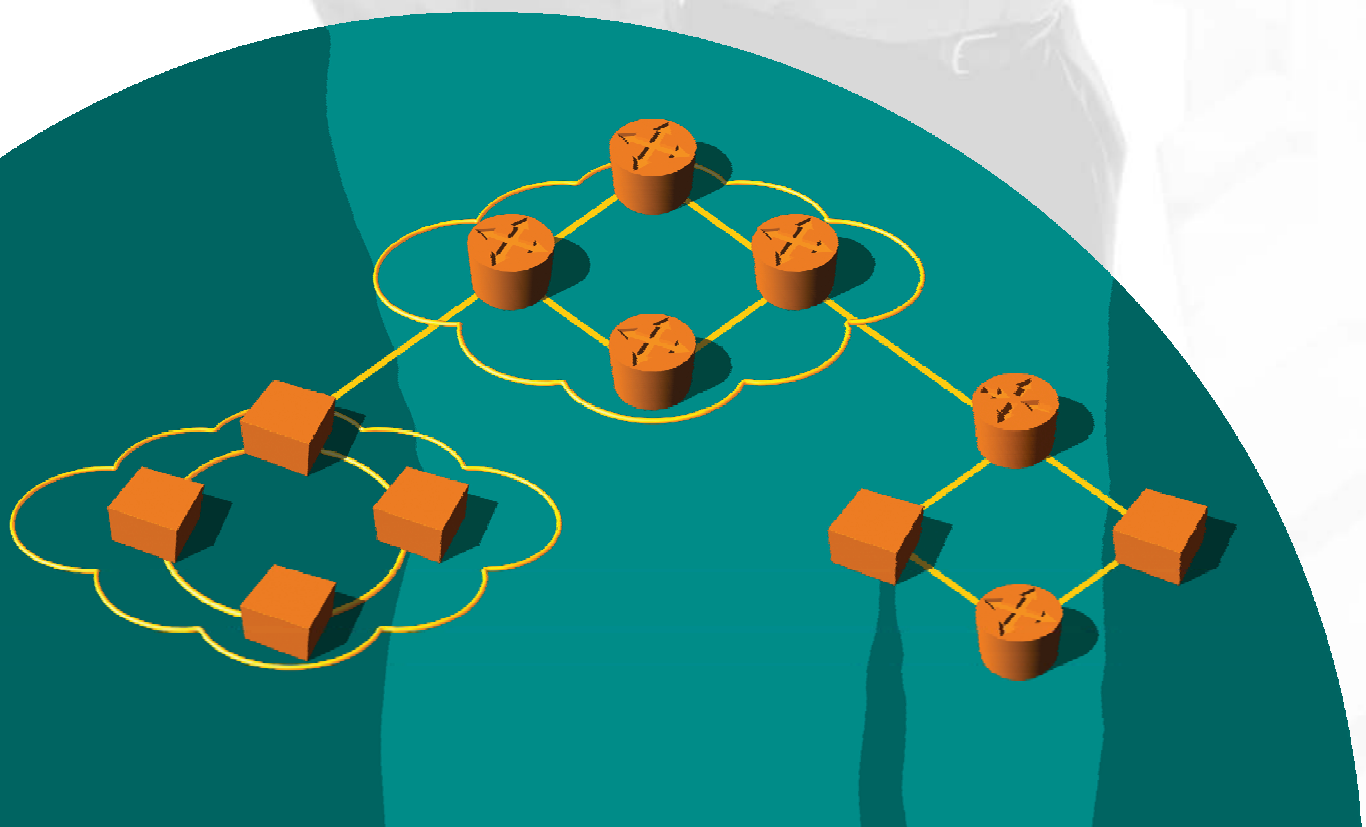




Analyze | Assure | Accelerate™

Spirent Communications Test Methodologies Journal Edge Router Edition



Introduction

Welcome to the Edge Router Edition of Spirent Communications Test Methodologies, a resource book containing tests that are vital to service providers and equipment manufacturers.

Edge router performance is increasingly important for building service providers' networks, though these devices are no longer just "on-ramps" or aggregation points. By "moving complexity to the edge," network functionality and the added value of services are distributed among hundreds or thousands of edge routers instead of overwhelming the network's core.

This architecture distributes the processing burden of high-touch applications such as firewalls, VPNs or QoS, and it assures that services are applied close to users to improve the customer's experience. The edge routers have become the heart of a carrier's service and are being seriously taxed by user demand.

Edge routers have complex jobs. Consequently, they have equally complex testing requirements. They typically support several concurrent protocols while simultaneously manage thousands of users. Protocol conversion and encapsulation must occur seamlessly and instantaneously. The only way to validate this complex functionality is to subject the devices to a series of conformance, performance and functional tests.

The Edge Router Edition provides suggested test methodologies for edge routers. Popular protocols such as OSPF and BGP are included. Additionally, test methodologies for emerging protocols such as MPLS (RSVP-TE and LDP) and MPLS VPNs are also provided. And finally, basic IP performance and QoS test methodologies are offered for IPv4 and IPv6.

Additional test methodologies can be found at Spirent Communications Developers Network at <http://scdn.spirentcom.com>, where you can also check for automated scripts.

The trend in moving complexity to the edge requires a new viewpoint on testing. We would like to help you achieve this expertise.

Sincerely,

Spirent Communications

Table of Contents

Spirent Edge Router Test Methodologies can be found on the Internet at Spirent Communications Developers Network <<http://scdn.spirentcom.com>>.

■ IP Throughput and Latency for IPv4 and IPv6 (TM-0100)	3
■ Packet Loss and Latency for IPv4 and IPv6 (TM-0101)	5
■ Back to Back (Burst Size) for IPv4 and IPv6 (TM-0102)	7
■ IP Quality of Service (TM-0139).....	9
■ OSPF Route Reconvergence (TM-0110)	11
■ RSVP-TE Forwarding Performance for Ingress LER (TM-0115)	14
■ RSVP-TE Forwarding Performance for Egress LER (TM-0116)	16
■ LDP Forwarding Performance for Ingress LER (TM-0112)	18
■ LDP Forwarding Performance for Egress LER (TM-0113)	20
■ BGP/MPLS VPN VRF Scalability (TM-0118)	22
■ Layer 2 Frames Over MPLS Scalability (TM-0119)	24
■ Acronyms	29
■ Glossary	30
■ Spirent Test Methodologies Information	34

IP Throughput and Latency for IPv4 and IPv6

- RFC 1242: Benchmarking Terminology for Network Interconnection Devices
 - RFC 2544: Benchmarking Methodology for Network Interconnect Devices
-

Objective

This data plane test is designed to provide throughput and latency information for a single router (device under test or DUT) or for a system of interconnected routers (system under test or SUT).

The throughput of a device or system is the maximum packet-forwarding rate for which the device or system will not drop any of the offered packets. Any packet loss can induce significant delays in the execution of higher layer applications; thus, knowing the maximum data rate a device or system can support without any packet loss is of crucial importance when judging the performance of a router or system of interconnected routers.

This test also determines the latency of a device or system (the time it takes a packet to travel through the device or the system), calculated at the maximum forwarding rate for which no packet loss is experienced (throughput rate).

The test methodology is applicable to IPv4 and IPv6 configurations.

Overview

To determine the throughput and latency of a device under test (DUT) or system under test (SUT), a minimum of two test ports will be required. All ports will be connected to the DUT/SUT.

One or more test ports will act as data transmitters and will offer traffic to the DUT/SUT. The other ports will act as data receivers and will accept traffic from the DUT/SUT. The DUT/SUT must be configured such that traffic offered by the data transmitters will be forwarded toward the data receivers.

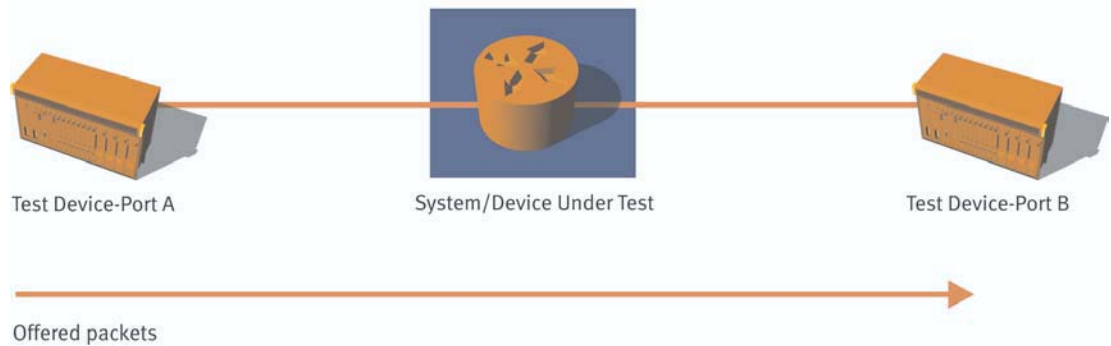
From the transmitter test ports, a predetermined number of packets is offered to the DUT/SUT. The packets are forwarded by the DUT/SUT to the receiver test ports. The number of packets accepted at the receiver test ports is compared to the number of packets transmitted. If packet loss occurs, then the offered load is decreased and the test is repeated. If packet loss is not observed, the test is repeated with an increased number of offered packets. By implementing a binary search pattern, the maximum rate for which no packet loss occurs is recorded. This rate corresponds to the DUT/SUT's throughput, or the first measurement of the test.

To calculate latency, a predefined test stream is delivered to the DUT/SUT from the source test ports at the calculated throughput rate. The transmitting timestamp, corresponding to when the test packet is sent by a transmitter test port, is subtracted from the receiving timestamp that corresponds to when the packet arrives at a receiver test port. The difference between the receiving timestamp and the transmitting timestamp gives the latency for a test packet. The test is run for a determined period of time, which should be long enough to collect sufficient data for an accurate representation of the system's latency.

Setup

In the baseline test, two test ports are used: Test Port Port A is the data transmitter, while Test Port B is the data receiver. For more complex tests, multiple data transmitters and multiple data receivers can be used.

IP Throughput and Latency Test for IPv4 and IPv6



Test Steps

1. Configure the DUT/SUT so the traffic offered from transmitter port(s) will be forwarded to the receiver port(s). This can be accomplished via a routing protocol or statically configured routes.
2. Configure the test parameters:
 - a. Initial packet rate.
 - b. Packet size.
 - c. Resolution rate for the binary search.
 - d. Latency test run time.
3. Send packets from source port(s):
Start with the maximum packet rate supported by the test port(s). Measure the number of packets offered at the transmitter port(s) and the number of packets accepted at the receiver port(s).
4. If no packet loss occurs, the throughput rate is equal to the offered data rate. Stop the test.
5. If packet loss occurs, decrease the packet rate and repeat the test.
6. Continue the binary search algorithm until the maximum packet rate with no packet loss is calculated (throughput rate).
7. Send a packet stream at the calculated throughput rate and measure the latency for the specified duration of time.

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Packet size.
- Resolution rate for the binary search.
- Latency test run time.

Test Outcome

Throughput, latency.

Packet Loss and Latency for IPv4 and IPv6

- RFC 1242: Benchmarking Terminology for Network Interconnection Devices
 - RFC 2544: Benchmarking Terminology for Network Interconnect Devices
-

Objective

This data plane test determines the packet loss rate and latency for a single router (device under test or DUT) or for a system of interconnected routers (system under test or SUT) for various input data rates and packet sizes.

The packet loss rate of a device or system is the percentage of Layer 3 frames that were offered at the input of the device or system but were not forwarded by the device or system due to hardware and software limitations. Calculating the packet loss rate of a system under different load conditions (input data rate and packet size) serves to evaluate how the system will perform under similar conditions in real-life operation.

This test also determines the latency of a device or system (the time it takes a packet to travel through the device or the system) calculated for the various input data rates, for which packet loss may or may not be experienced.

The test methodology is applicable to IPv4 and IPv6 configurations.

Overview

To determine the packet loss rate and latency of a DUT or SUT, a minimum of two test ports will be required. All ports will be connected to the DUT/SUT.

One or more test ports will act as data transmitters and will offer traffic to the DUT/SUT. The other ports will act as data receivers and will accept traffic from the DUT/SUT. The DUT/SUT must be configured such that the traffic offered by the data transmitters will be forwarded toward the data receivers.

From the data transmitters, a predetermined number of packets is offered to the DUT/SUT for a given amount of time. The packets are forwarded by the DUT/SUT to the receiver test ports. The number of packets accepted at the receiver test ports (RxPacketCount) is compared to the number of packets offered from the transmitter ports (TxPacketCount).

The packet loss rate is calculated using the following formula:

$$\frac{[(TxPacketCount - RxPacketCount) \times 100]}{TxPacketCount}$$

The packet loss is calculated for input data rates starting with 100 percent of the maximum rate that can be offered from the transmitter test ports. The input data rate is decremented and the test repeated, until there are two successive trials where there is no packet loss. The test may be stopped when the input data rate reaches a user-selected threshold beyond which no measurements are required. The amount by which the input data rate is decremented should be at most 10 percent of the maximum input data rate and can be as low as 1 percent.

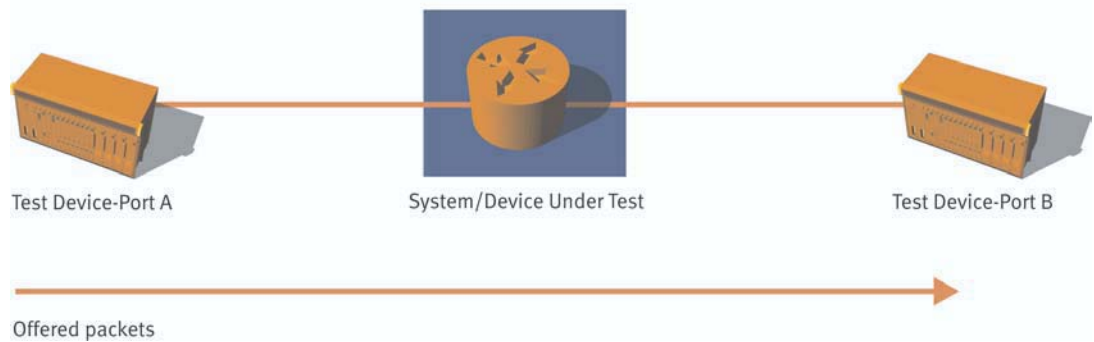
System latency is measured at every trial; thus, for every value of the input data rate.

To calculate latency, a predefined test stream is delivered to the DUT/SUT from the transmitter test ports at the calculated throughput rate. The transmitting timestamp, corresponding to when a test packet is sent by a transmitter test port, is subtracted from the receiving timestamp that corresponds to when the packet arrives at a receiver test port. The difference between the receiving timestamp and the transmitting timestamp give the latency for a test packet. The test is run for a determined period of time, which should be long enough to collect sufficient data for an accurate representation of the system's latency.

Setup

The configuration of the test is presented in the setup diagram. In the baseline test, two test ports are used: Test Port A is the data transmitter, while Test Port B is the data receiver. For more complex tests, multiple data transmitters and multiple data receivers can be used.

Packet Loss and Latency Test for IPv4 and IPv6



Test Steps

1. Configure the DUT/SUT so the traffic offered from the transmitter port(s) is forwarded to the receiver port(s). This can be accomplished via a routing protocol or statically configured routes.
2. Configure the test parameters:
 - a. Initial packet rate.
 - b. Packet size.
 - c. Resolution by which the input data rate is decremented.
 - d. Run time for every trial.
3. Send packets from transmitter port(s):
Start with the maximum packet rate supported by the test port(s). Measure the number of packets offered at the transmitter port(s) and the number of packets accepted at the receiver port(s).
4. If no packet loss occurs in two successive trials, stop the test.
5. If packet loss occurs, calculate the packet loss rate and the latency.
6. Decrement the input data rate and repeat test.
7. The test stops when no packet loss is recorded in two successive trials, or when the input data rate reaches a user-selected threshold.

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Packet size.
- Test time.
- Packet rate.
- Rate increment/decrement resolution.

Test Outcome

Packet loss, latency.

Back-to-Back (Burst Size) Test for IPv4 and IPv6

- RFC 1242: Benchmarking Terminology for Network Interconnection Devices
 - RFC 2544: Benchmarking Terminology for Network Interconnect Devices
-

Objective

This test determines the maximum number of packets a router (device under test or DUT), or system of interconnected routers (system under test or SUT), can forward “back-to-back” without packet loss. The number of packets in the longest burst that does not cause packet loss is the back-to-back value.

In a back-to-back test, packets are delivered at full line rate with no pause between successive packets, except the required “legal” separation for a given technology and physical medium.

The test methodology is applicable to IPv4 and IPv6 configurations.

Overview

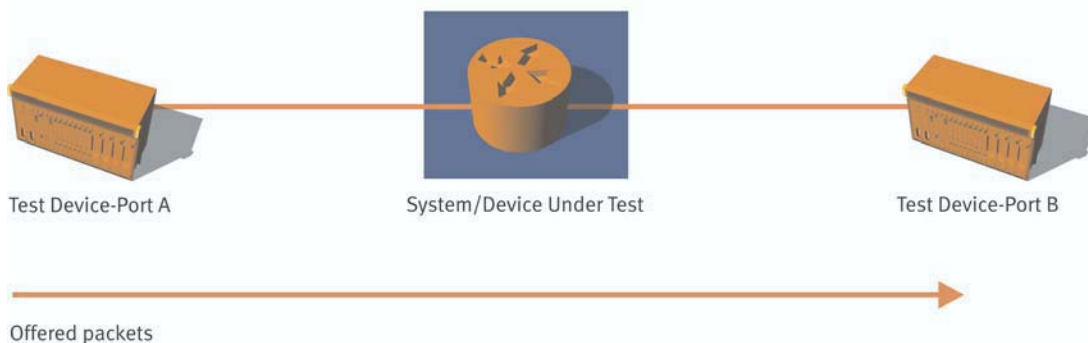
To determine the back-to-back value of a device under test (DUT) or system under test (SUT), a minimum of two test ports will be required. All ports will be connected to the DUT/SUT.

One or more test ports will act as data transmitters and offer traffic bursts to the DUT/SUT. The other ports will act as data receivers and will accept traffic from the DUT/SUT. The DUT/SUT must be configured such that the traffic offered by the data transmitters will be forwarded toward the data receivers.

From the transmitter test ports, a burst of back-to-back packets is offered to the DUT/SUT. The packets are forwarded by the DUT/SUT to the receiver test ports. The number of packets accepted at the receiver test ports is compared to the number of packets offered. If packet loss occurs, then the burst size is decreased (the number of back-to-back packets is decreased) and the test is repeated. If no packet loss is observed in the next iteration, then the burst size is increased and the test is repeated. By implementing a binary search, the back-to-back value is determined.

Setup

Back-to-Back Test for IPv4 and IPv6 Capable Routers



In the baseline test, two test ports are used: Test Port A is the data transmitter, while Test Port B is the data receiver. For more complex tests, multiple data transmitters and multiple data receivers can be used.

Test Steps

1. Configure the DUT/SUT so traffic offered from transmitter port(s) is forwarded to the receiver port(s). This can be accomplished via a routing protocol or statically configured routes.
2. Configure the test parameters:
 - a. Packet size.
 - b. Burst size (number of packets).
 - c. Packet rate.
 - d. Burst increment/decrement resolution (for the binary search).
3. Send packets from transmitter port(s) for a determined period of time. Start at the full line rate; measure the number of packets offered at the transmitter port(s) and the number of packets accepted at the receiver port(s).
4. If no packet loss occurs, the DUT/SUT can handle back-to-back packets coming at full line rate; thus, there is no burst size limitation. Stop the test.
5. If packet loss occurs, decrease the burst size and repeat the test.
6. Continue the binary search algorithm until the maximum burst size with no packet loss is determined (the back-to-back value).

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Packet size.
- Burst size.
- Packet rate.
- Burst increment/decrement resolution.

Test Outcome

Back-to-back value.

IP Quality of Service

- RFC 1349: Type of Service in the Internet Protocol
 - RFC 2474: Definition of the DS Field
 - RFC 2475: Differentiated Services Architecture Model
-

Objective

This test validates a router's ability to correctly prioritize, queue and forward IP datagrams. Though this methodology is commonly referred to as a quality of service (QoS) test, it should more accurately be called a class of service (CoS) or IP prioritization test.

This test will be applicable to routers that support IP service prioritization based on the Differentiated Service (DiffServe) code points and the IP type of service (ToS) field.

For the purposes of this document, the test methodology will be based on IPv4, but it also is equally applicable to IPv6 data by using the "Traffic Class" field in the packet header.

Overview

This is a data plane performance and functional test. It will require two ports (ingress and egress) on the router and two corresponding ports on the test equipment. The most dramatic results can be obtained if the ingress port supports considerably higher bandwidth than the egress port. For example, use a Gigabit Ethernet port for traffic generation, and then use a 100 Mbps Ethernet port as the destination.

From the data source test port, generate three streams of IP traffic, each with different Differentiated Services or ToS characteristics. All of these streams should be destined for simulated hosts on the egress port. For example, if the router's egress port is an Ethernet port with a local address of 192.1.1.1/24, configure the data traffic to have destination addresses of 192.1.1.2, 192.1.1.3 and 192.1.1.4. It is recommended (though not necessary) that each stream uses a different destination address. This will simplify viewing and analyzing of results.

The router (DUT) should be enabled to support IP QoS traffic shaping on the egress port. This can be implemented in two different ways:

The first option uses a strict queuing mechanism, in which all high priority traffic will be transmitted prior to servicing the medium or low priority queues. If this method is used, it should be easy to generate a "queue starvation" scenario in which no low priority traffic is transmitted.

The alternative is based on a weighted fair queuing algorithm. This procedure assigns a weighted value to each of the traffic queues. For example, high priority traffic might be serviced 60 percent of the time; medium priority 30 percent of the time and low priority traffic will be entitled to 10 percent of the bandwidth. Regardless of which method is implemented, it should be easy to validate the results as the traffic is routed through the DUT.

Multiple iterations of this test should be conducted. The load factors and packet sizes should be varied to validate the router's performance under different network conditions.

There currently are two competing models for IP class of service implementations, both of which are dependent upon the ToS Field in the IP packet's header. The first option is specified by RFC 1349 and defines the expected characteristics for specific ToS values. The other is the DiffServe model specified by RFCs 2474 and 2475. This is the more common option. It is critical that users understand which of these models their router vendor implements.


Setup

The physical setup for this particular test is quite simple, as indicated in the diagram below:

In this diagram, the data traffic will be generated at Port 1 (the ingress to the router); the results will be analyzed at Port 2 (the egress).

Test Steps

1. Configure IP addresses on two router ports. For this example, the egress port will be 192.1.1.1/24. If necessary, build a static route between these two ports.
2. Configure the QoS parameters for three different service classes for the egress port.
3. Using the traffic generator connected to Port 1 of the router, construct three different IP traffic streams destined to be simulated hosts on the 192.1.1.0/24 network. Each traffic stream will have a different ToS value. These values must correspond to the service classes defined on the router. It is suggested the aggregate input traffic oversubscribe the egress bandwidth.
4. Verify packets are being correctly forwarded to the egress port.
5. Verify the router is correctly prioritizing (and even dropping) packets based on the predetermined ToS values.
6. If the router is capable of modifying the ToS values, this function can also be verified in a subsequent test.



Test Parameters

It is highly likely that different outcomes will be obtained by modifying any of the following input parameters:

- Packet size.
- Quantities of input streams.
- Packet rate.
- Distribution of the input streams.

Test Outcome

Pass/Fail and the limits associated with the router's implementation of IP quality of service.

OSPF Route Reconvergence

- RFC 2328: OSPF Version 2
 - Draft-bmwg-ospfconv-term-00.txt: OSPF Convergence Testing Terminology and Concepts
 - Draft-bmwg-ospfconv-intraarea-00.txt: Benchmarking Methodology for Basic OSPF Convergence
-

Objective

This test determines the time for an OSPF router or system of interconnected OSPF routers to reconverge route tables and redirect data traffic when the best route to a destination is unavailable. This process is also called reconvergence, as the router “reconverges” its route tables to the new topology. Measuring the reconvergence performance of OSPF routers is important to determine the IP network speed recovery and data loss when confronted with failures or changes in the topology.

Overview

To determine the route reconvergence performance of a device under test (DUT) or system under test (SUT), a minimum of three test ports will be required. All ports will be connected to the DUT/SUT.

Two or more test ports will advertise to the DUT/SUT possible routes to the same destination nodes or hosts, but only one of the test ports will have the lowest cost to the destination. This test port will be selected by the DUT as the best path towards the destination.

After the best path is installed in the DUT/SUT, all data traffic to the given destination will be carried on the best path.

Another test port acting as a data source will send traffic to the given destination. The data traffic will reach the DUT/SUT, which will use the best path to forward the data. Thus, all data is directed to the test port that has the lowest cost to the destination network behind it.

As part of the test, the best path will become unavailable. Thus the test port situated on the best Path will generate Link State Update messages to the DUT advertising the path’s unavailability. An alternate path to the same destination is selected by the DUT to forward the data. After this selection process is completed, the traffic is switched to the alternate path. The time it takes the router to complete the reconvergence process is evaluated.

Several measurements related to the control and data planes can be obtained:

RouterWithdrawTime is the time it takes a router to completely withdraw the best path from its tables once it receives an LSA announcing the path’s unavailability. *RouterWithdrawTime* is calculated as the time between the moment the first withdrawing Link State Update message is sent and the moment the last data packet is received on the best path.

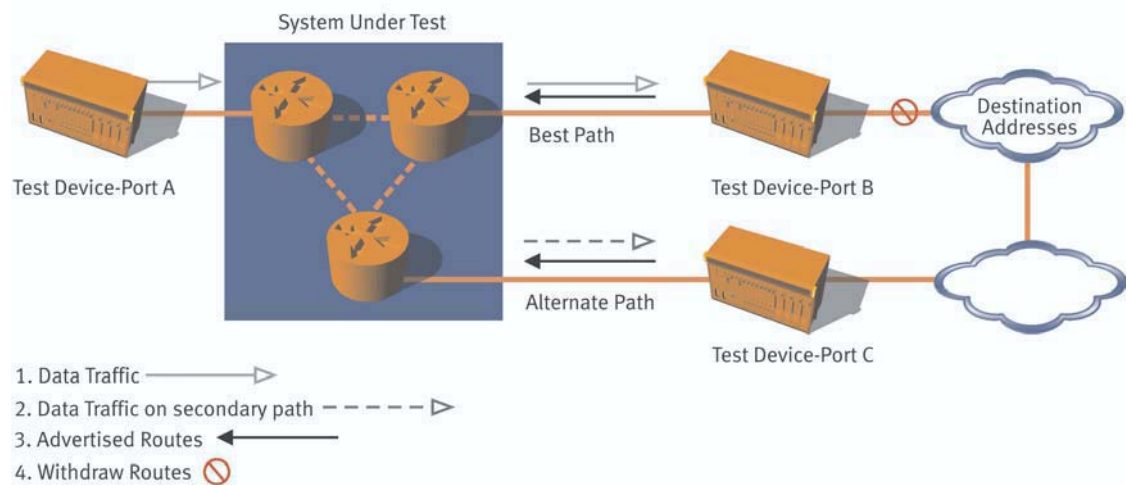
RouterSwitchTime is the time it takes a router to switch all packets from the best path to the alternate path. *RouterSwitchTime* can be calculated as the time between the moment the first data packet is received on the alternate path and the moment the traffic level on the alternate path has reached the same level as it was on the best path.

TotalConvergenceTime is the time it takes a router to completely switch the traffic onto the alternate path once it receives an LSA announcing the path’s unavailability.

These measurements characterize a router’s overall reconvergence performance, as the test involves operations related to the control plane as well as data plane.

Setup

OSPF Route Reconvergence Time Test



For this evaluation, three of the ports on the DUT/SUT must be connected to test devices for the baseline convergence test. More than three ports can be used for more complex tests.

A, B and C represent ports on the test device(s) that is (are) connected to the DUT/SUT.

Port A is the traffic source from which data packets are sent to the destination addresses.

Port B is a test port that advertises the best route to the destination addresses.

Port C is a test port that advertises the alternate route to the destination addresses.

Test Steps

1. Create an emulated topology behind ports B and C.
2. Establish OSPF adjacencies between the DUT/SUT and ports B and C.
3. Send LSAs from Port B and Port C to the DUT/SUT advertising connectivity to a set of destination addresses. The route going through Port B will have a lower cost than the route going through Port C. Thus, the route going through Port B will be selected as the best path to the set of destination routes and will be included in the Shortest Path Tree calculated from the DUT/SUT.
4. Send traffic from the source test port (Port A) to the destination addresses that were previously advertised by ports B and C.
5. Verify that traffic is received at Port B.
6. Send LSAs from the test port associated with the best path (Port B) withdrawing the routes to the destination addresses. Record the time the corresponding Link State Update message was sent from Port B (to be used when calculating *RouterWithdrawTime* and *TotalConvergenceTime*). Record the timestamp of the last packet received on the best path by Port B (also to be used when calculating *RouterWithdrawTime*).
7. Verify that the data traffic is being switched from the best path to the alternate path. Record the timestamp of the first data packet received on the alternate path by Port C (to be used when calculating *RouterSwitchTime*). Record the timestamp corresponding to the moment the traffic level on the alternate path has reached the same level as it was on the best path (to be used when calculating *RouterSwitchTime* and *TotalConvergenceTime*).
8. Calculate *RouterWithdrawTime*, *RouterSwitchTime*, *TotalConvergenceTime* and eventual packet loss.

Test Parameters

Different test outcomes can be obtained by modifying any of the following input parameters:

- Total number of destination addresses advertised by ports B and C.
- Load offered from the traffic source (from Port A).
- Length of packets sent from the traffic source (from Port A).

Test Outcome

RouterWithdrawTime, RouterSwitchTime, TotalConvergenceTime, packet loss.

RSVP-TE Forwarding Performance for Ingress LER

- RFC 2205: Resource ReSerVation Protocol
- RFC 3209: “RSVP-TE: Extensions to RSVP for LSP Tunnels”

Objective

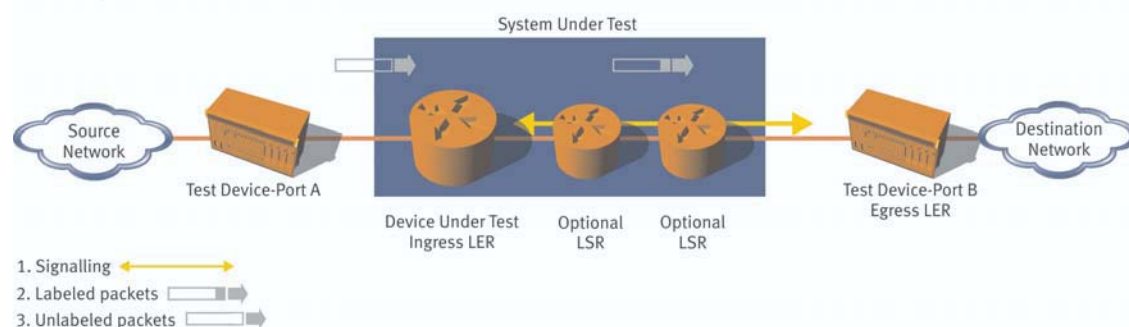
This test verifies the device under test (DUT) is capable of establishing RSVP-TE LSP tunnels as an Ingress LER and correctly labeling and forwarding MPLS labeled packets.

Overview

LSPs are created with the DUT as the Ingress LER and Test Port B as the Egress LER. The source (Test Port A) will supply unlabeled packets to the DUT with addresses corresponding to the LSPs connecting the ingress and egress nodes. The DUT will forward MPLS labeled packets to the Egress LER, which will capture the forwarded labeled packets for verification. The system under test (SUT) may consist of a single DUT or, optionally, a DUT with additional routers. If optional routers are used so that the DUT and test ports are not directly connected, a route distribution method (OSPF, IS-IS, static routes) may be necessary.

Setup

RSVP-TE Forwarding Performance Test
DUT = Ingress



Test Steps

1. If the DUT is not directly connected to the test ports, implement a route distribution method (OSPF, IS-IS, static routes) to advertise the address of the Egress LER to the Ingress LER.
2. Configure DUT to support MPLS RSVP-TE and setup a static LSP with the Egress LER as the tunnel destination.
3. Send unlabeled packets from the source test Port A to an address destination serviced by the LSP.
4. Analyze MPLS labeled packets received and Egress LER.
5. If unlabeled packets are received:
 - a. Record traffic statistics and a Fail verdict.
 - b. End the test.

6. If the number of labeled packets received is equal to the number of unlabeled packets sent:
 - a. Record traffic statistics and a Pass verdict.
7. To characterize router behavior, vary one of the following and continue at Step 3. Otherwise, end the test.
 - a. Packet length.
 - b. Offered load.

Test Parameters

- Packet length.
- Offered load.

Test Outcome

Traffic and RSVP-TE statistics.

RSVP-TE Forwarding Performance for Egress LER

- RFC 2205: Resource ReSerVation Protocol
- RFC 3209: “RSVP-TE: Extensions to RSVP for LSP Tunnels”

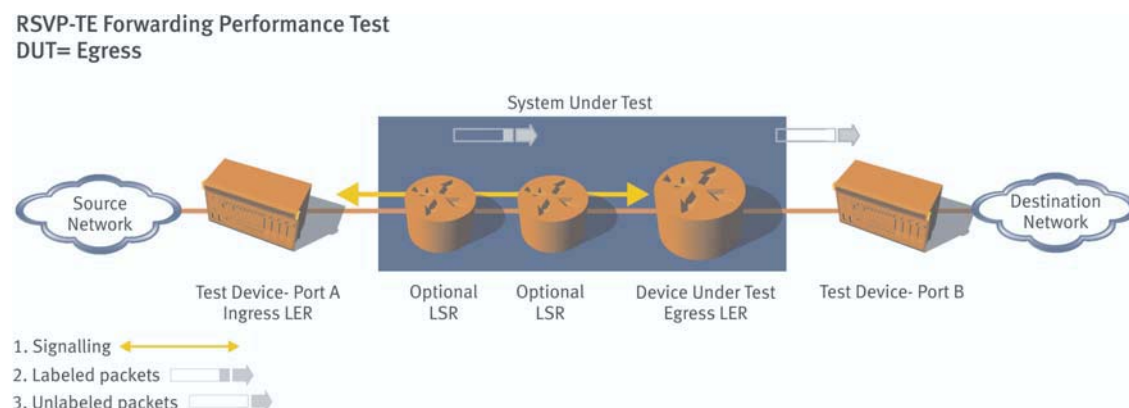
Objective

This test verifies the device under test (DUT) is capable of creating RSVP LSP tunnels as an Egress LER, and to correctly remove the MPLS label from received packets and forward unlabeled packets to the destination router.

Overview

LSPs are created with Test Port A as an Ingress LER and the DUT as the Egress LER. The Ingress LER will supply MPLS labeled packets to the DUT using a label allocated by the DUT. The destination (Test Port B) will receive the unlabeled packets forwarded by the DUT and provide statistics. The offered load and packet length may be varied to more thoroughly characterize router behavior. The system under test (SUT) may consist of a single DUT or, optionally, a DUT with additional routers. If optional routers are used so the DUT and test ports are not directly connected, a route distribution method (OSPF, IS-IS, static routes) may be necessary.

Setup



Test Steps

1. If the DUT is not directly connected to the test ports, implement a route distribution method (OSPF, IS-IS, static routes) to advertise the address of the Egress LER to the Ingress LER.
2. Send a Path message from the Ingress LER to setup an LSP with the DUT as the tunnel end point.
3. Send MPLS labeled packets from the Ingress LER using the label that is allocated from the DUT. The destination router of the IP packet is the test port directly attached to the DUT.
4. Analyze unlabeled packets received at destination Test Port B.

5. If labeled packets are received:
 - a. Record traffic statistics and a Fail verdict.
 - b. End the test.
6. If the number of unlabeled packets received is equal to the number of labeled packets sent, record traffic statistics and a Pass verdict.
7. To characterize router behavior, vary one of the following and return to Step 3. Otherwise, end the test.
 - a. Packet length.
 - b. Offered load.

Test Parameters

- Packet length.
- Offered load.

Test Outcome

Traffic and RSVP-TE statistics.

LDP Forwarding Performance for Ingress LER

- RFC 3036: LDP Specification
- RFC 2328: OSPF version 2

Objective

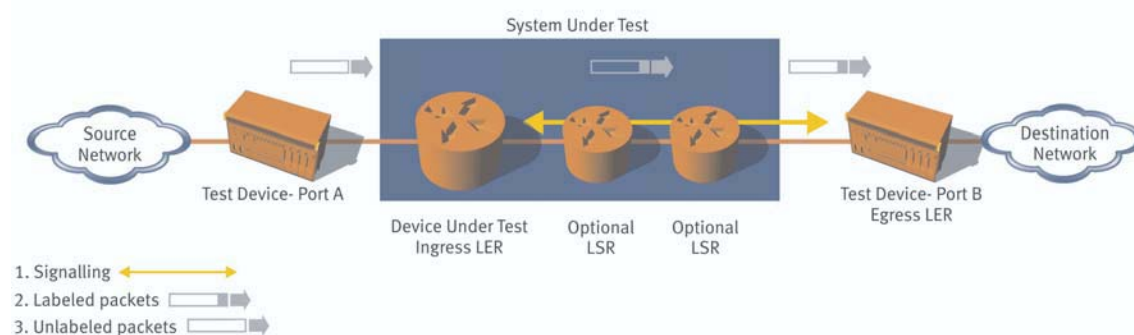
This test verifies the device under test (DUT) is capable of creating LSP Tunnels as an Ingress LER using LDP, correctly attaching the label to received unlabeled packets and forwarding MPLS labeled packets to the Egress LER.

Overview

An LSP is created with the DUT as the Ingress LER and Test Port B as the Egress LER. The source (Test Port A) will supply unlabeled data packets to the DUT. The DUT will attach the appropriate label to the packet and forward it through the LSP to the Egress LER. The Egress LER test port will receive the forwarded MPLS labeled packets and provide statistics. The number of LSPs and/or data rate may be increased at each iteration until packet loss is observed. The system under test (SUT) may consist of a single DUT or, optionally, a DUT with additional routers. If optional routers are used so that the DUT and test ports are not directly connected, a route distribution method (OSPF, IS-IS, static routes) may be necessary.

Setup

(CR) LDP Forwarding Performance Test
DUT= Ingress LER



Test Steps

1. If the DUT is not directly connected to the test ports, implement a route distribution method (OSPF, IS-IS, static routes).
2. Advertise a set number of destination routes from Test Port B.
3. Configure LDP label distribution method on the DUT and appropriate test ports as either Downstream on Demand (DoD) or Downstream Unsolicited (DU).
4. Depending on the method of label distribution (DoD, DU), the MPLS labels will be distributed by the downstream node (DU) or the labels will be requested by upstream nodes (DoD).
5. Send unlabeled data packets from Test Port A to all advertised destinations at a set rate.

6. Analyze MPLS labeled packets received at Egress LER.
7. If packet loss is observed:
 - a. Record traffic statistics and a Fail verdict.
 - b. End the test.
8. If no packet loss is observed:
 - a. Record traffic statistics and a Pass verdict.
9. To characterize router behavior, vary one of the following and return to Step 3. Otherwise, end the test.
 - a. Number of LSPs.
 - b. Packet length.
 - c. Packet rate.

Test Parameters

- Number of LSPs.
- Packet length.
- Packet rate.

Test Outcome

Traffic and LDP statistics.

LDP Forwarding Performance for Egress LER

- RFC 3036: LDP Specification
- RFC 2328: OSPF version 2

Objective

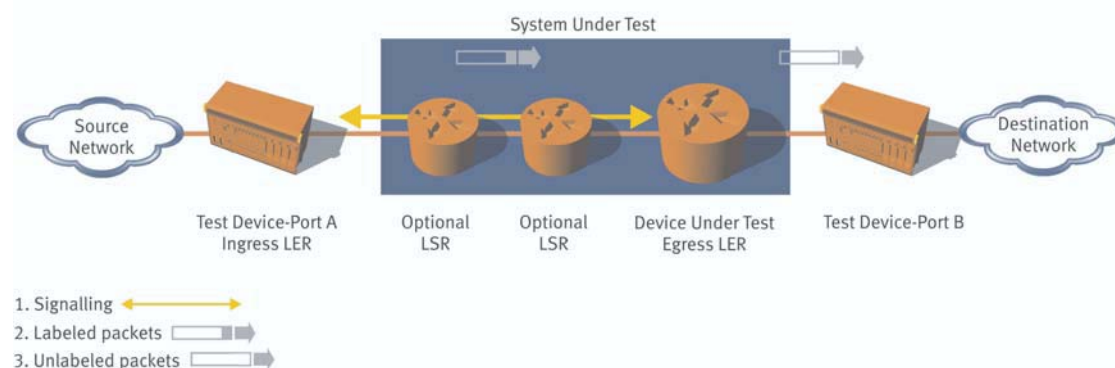
This test verifies the device under test (DUT) is capable of creating LSP Tunnels as an Egress LER using LDP to distribute labels, correctly removing the label from received labeled packets and forwarding unlabeled IP data packets to the destination router.

Overview

An LSP is created with Test Port A as an Ingress LER and the DUT as the Egress LER. The Ingress LER will supply MPLS labeled packets to the DUT using a label distributed by the DUT. The destination test port will receive the unlabeled packets forwarded by the DUT and provide statistics. The system under test (SUT) may consist of a single DUT or, optionally, a DUT with additional routers. If optional routers are used so that the DUT and test ports are not directly connected, a route distribution method (OSPF, IS-IS, static routes) may be necessary.

Setup

(CR) LDP Forwarding Performance Test
DUT = Egress LER



Test Steps

1. If the DUT is not directly connected to the test ports, implement a route distribution method (OSPF, IS-IS, static routes).
2. Advertise a set number of destination routes from Test Port B.
3. Configure LDP label distribution method on the DUT and appropriate test ports as either Downstream on Demand (DoD) or Downstream Unsolicited (DU).
4. Depending on the method of label distribution (DoD, DU), the MPLS labels will be distributed by the downstream node (DU) or the labels will be requested by upstream nodes (DoD).
5. Send labeled data packets from the Ingress LER to all advertised destinations at a set rate.

6. Analyze IP packets received at the destination Port B.
7. If packet loss is observed:
 - a. Record traffic statistics and a Fail verdict.
 - b. End the test.
8. If no packet loss is observed, record traffic statistics and a Pass verdict.
9. To characterize router behavior, vary one of the following and return to Step 3. Otherwise, end the test.
 - a. Number of LSPs.
 - b. Packet length.
 - c. Packet rate.

Test Parameters

- Number of LSPs.
- Packet length.
- Packet rate.

Test Outcome

Traffic and LDP statistics.

BGP/MPLS VPN VRF Scalability

- Draft-ietf-ppvpn-rfc2547bis-xx: BGP/MPLS VPNs
- RFC 2858: Multi-Protocol Extensions for BGP-4
- Draft-ramachandra-bgp-ext-communities-xx: BGP Extended Communities Attribute
- RFC 3107: Carrying Label Information in BGP-4

Objective

This test will determine the number of VPN routing and forwarding tables (VRFs) and routes per VRF a provider edge (PE) router in a BGP/MPLS VPN system can support.

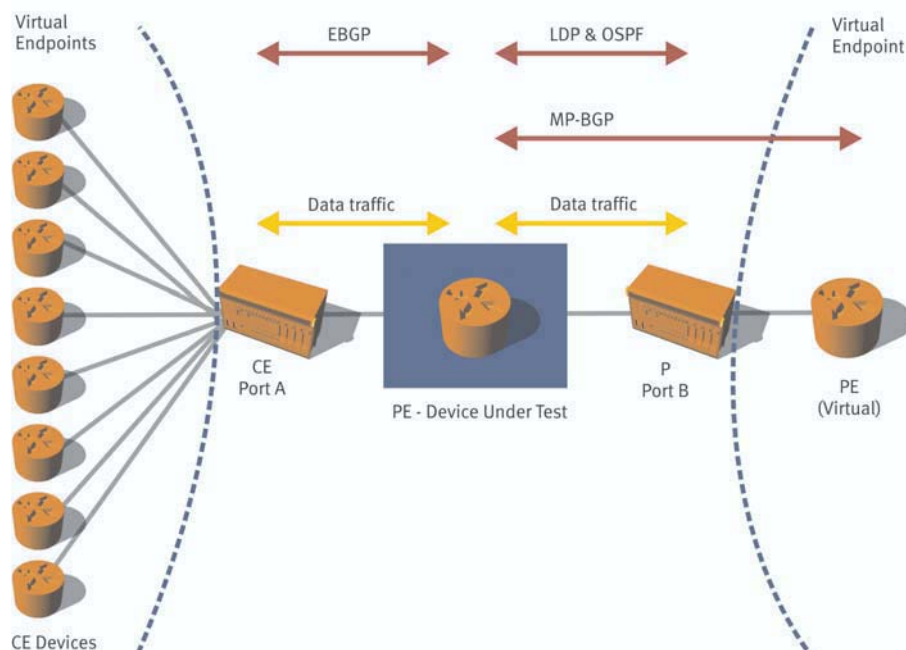
Overview

BGP/MPLS VPNs use Layer 3 routing information to distribute VPN connectivity information. An MPLS backbone network provides data plane connectivity. Customer edge (CE) routers advertise their site reachability to their directly connected provider edge (PE) routers via a routing protocol. The PE router installs these routes into a VRF for that VPN. These VPN routes are then propagated to other PE routers to provide this reachability information to remote VPN sites.

This test uses eBGP to distribute routes between CE and PE (DUT). Some type of sub-interfacing is used to easily scale the test to many customer sites. Sub-interfacing will create a correspondence between multiple logical interfaces and a single physical interface. This is typically accomplished by assigning to each logical interface layer 2 channel identifiers, such as Ethernet VLAN IDs, Frame Relay DLCI or ATM VPI/VCI. The DUT should propagate these VPN routes via MP-BGP to the PE test port where they can be verified. LDP is used as the core MPLS protocol. Data traffic will be used to test the validity of VPN routes.

Setup

BGP/MPLS VPN VRF Scalability Test



Test Steps

1. Configure a set number of sub-interfaces on each interface used in the test between the CE and PE router (DUT).
2. Provision and associate a different VRF for each configured sub-interface.
 - a. Configure the VRF on DUT.
 - b. Configure the Route Distinguishers.
 - c. Configure the import and export route targets. This will help to ensure routes are installed in the proper VRFs.
 - d. Configure the association between the sub-interfaces and the VRFs.
 - e. Configure Multi-protocol BGP on the core-facing interface of the DUT.
 - f. Configure OSPF and LDP on the core-facing interface of the DUT.
3. Configure eBGP between the CE-PE interfaces on each sub-interface.
4. Advertise a set number of routes from each CE to the PE.
5. Monitor the propagated routes via the MP-BGP session. This verifies that all the routes advertised by the CEs are re-advertised properly.
6. If the propagated routes do not equal the routes advertised:
 - a. Record the number of VRFs and number of routes advertised and a Fail verdict.
 - b. End the test.
7. If the propagated routes equal the routes advertised:
 - a. Record the number of VRFs and number of routes advertised.
 - b. Use the advertised label mappings and prefixes to send data traffic from the core-facing test port to the CE test ports.
 - c. Record a pass verdict for VPNs that have properly forwarded packets. Record a Fail verdict for VPNs with packet forwarding errors.
 - d. End the test.
8. To characterize router behavior, vary the following and return to Step 1. Otherwise, end the test.
 - a. Number of physical interfaces (ports).
 - b. Number of routes advertised.

Test Parameters

- Number of physical interfaces.
- Number of CEs per physical interface.
- Number of routes advertised per CE.
- Offered load.
- Datagram length.

Test Outcome

Number of CEs (VRFs) created, number of routes advertised, number and list of valid VPNs, number and list of invalid VPNs, number of physical interfaces.

Layer 2 Frames Over MPLS Scalability

- Draft-martini-l2circuit-encap-mpls-xx
 - Draft-martini-ethernet-encap-mpls-xx
 - Draft-martini-frame-encap-mpls-xx
 - Draft-martini-atm-encap-mpls-xx
-

Objective

The goal is to test the functional behavior of a PE device when provisioning Martini virtual circuits and forwarding traffic over those established virtual circuits.

This test allows the operator to verify the proper control plane signaling and data plane forwarding of Martini-encapsulated traffic by checking the LDP extensions for Martini signaling, as well as the complex encapsulation and de-encapsulation of Martini traffic.

Overview

Two test ports are required for determining the Martini virtual circuit (VC) capacity of a system under test (SUT). All test ports will be connected to the SUT. One test port will be connected to the customer-facing side of the SUT, and one port will be connected to the core-facing side.

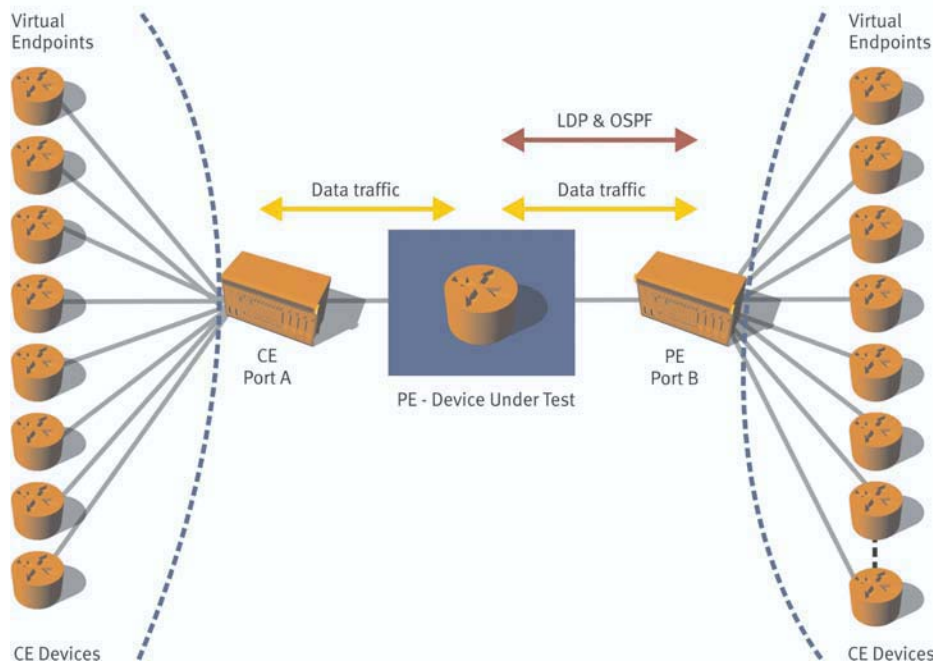
A user-specified number and type of VC is established via LDP on the core-facing port of the SUT. These circuits are mapped by the SUT to and from the customer-facing port through the core-facing tunnels. The label and virtual circuit information exchanged with the SUT's core-facing port is gathered for use in the generation and analysis of data traffic.

Layer 2 traffic is configured and offered to the customer-facing side of the SUT to be analyzed as Martini-encapsulated frames on the core-facing port of the SUT. Conversely, Martini-encapsulated traffic is configured and offered to the core-facing side of the SUT to be analyzed as Layer 2 traffic on the customer-facing side.

The number of valid VCs can be evaluated. Encapsulation/de-encapsulation problems can be reported.

Setup

Layer 2 Frames over MPLS



The test configuration and setup diagram, showing details such as test ports and traffic distribution, helps explain the test procedure.

The virtual circuit tunnels occur between the DUT (PE) and the right side test port (PE). The test creates virtual circuits to map the virtual customer endpoints on the left, with the virtual customer endpoints on the right.

Test Steps

1. Physical connectivity – This test requires at least two test ports. Connect at least one port to the “customer-facing” side of the DUT and one port to the “core-facing” side of the DUT.
2. Virtual circuit provisioning – Configure the DUT for the number of VCs needed for the test:
 - a. Configure the DUT for LDP and OSPF on the core-facing interface.
 - b. Map a Layer 2 circuit from a customer side interface to a VC ID and the PE test port IP address. The Layer 2 circuit will be identified by a VLAN ID, Frame Relay DLCI or ATM VPI/VCI.
 - c. Repeat Step 2b for all VCs needed.
3. Core facing test port configuration (PE test port):
 - a. Configure OSPF to establish adjacency with DUT.
 - b. Configure LDP to establish session with DUT.
 - c. Advertise interface address through LDP.
 - d. Configure VC FEC label mappings using Downstream Unsolicited (DU) mode for the VCs configured in Step 2. The same VC IDs and traffic type must be used. This is for traffic flowing in the opposite direction.
4. Traffic configuration.
 - a. Configure the customer-facing test ports (CE test ports) for generation
 - i. Define test streams for the Layer 2 circuits defined in Step 2b.
 - ii. Set the offered load and packet length parameters to the desired settings.

- b. Configure the customer-facing test ports for analysis – set the analyzer filter to the appropriate Layer 2 ID (VLAN ID, DLCI or VPI/VCI).
 - c. Configure the core-facing test ports for generation.
 - i. Define test streams by mapping the inner label values assigned by the DUT to the appropriate Martini encapsulation type and Layer 2 ID for each VC. Add the outer label for the core tunnel if applicable.
 - ii. Set the offered load and packet length parameters to the desired settings.
 - d. Configure the core-facing test ports for analysis set the analyzer filter to filter on the inner label value and appropriate Layer 2 ID.
5. Run test.

Test Parameters

- Number of configured VCs.
- Layer 2 ID and VC ID for each configured VC.
- Offered load, packet length and other traffic configuration parameters.

Test Outcome

1. VC statistics.
 - a. Number of VCs successfully established.
 - b. Encapsulation/de-encapsulation errors.
 - i. Packets switched with wrong Layer 2 ID on customer-facing interfaces.
 - ii. Packets switched with wrong inner label on core-facing interfaces.
 - c. Rx packet rate, packet loss, packet sequence errors, packet count.

Appendix

Acronyms

AS – Autonomous System
ATM – Asynchronous Transfer Mode
BGP – Border Gateway Protocol
CE – Customer Edge
DLCI – Data Link Connection Identifier
DoD – Downstream on Demand
DU – Downstream Unsolicited
DUT – Device Under Test
EBGP – Exterior Border Gateway Protocol
FEC – Forwarding Equivalence Class
FIB – Forwarding Information Base
ID – Identifier
IGP – Interior Gateway Protocol
IP – Internet Protocol
IPv4 – Internet Protocol Version 4
IPv6 – Internet Protocol Version 6
IS-IS – Intermediate System to Intermediate System
LDP – Label Distribution Protocol
LER – Label Edge Router
LSA – Link State Advertisements
LSP – Label Switched Path
MP-BGP – Multi-Protocol Border Gateway Protocol
MPLS – Multi-Protocol Label Switching
OSPF – Open Shortest Path First
PE – Provider Edge
RFC – Request For Comment
RIB – Routing Information Base
RSVP – Resource Reservation Protocol
RSVP-TE – Resource Reservation Protocol with Traffic Engineering
Rx – Receive
SUT – System Under Test
Tx – Transmit
VC – Virtual Circuit
VLAN – Virtual LAN
VPN – Virtual Private Network
VRF – VPN Routing and Forwarding

Glossary

A

Adjacency A relationship that is established between two routers in the process of exchanging routing protocol messages (example: BGP or OSPF adjacency).

AS (Autonomous System) A part of a network under a single administrative domain. Usually running a single internal routing protocol.

B

BGP (Border Gateway Protocol) The exterior gateway protocol used for distributing routes over the Internet. Currently, Version 4 (BGP-4) is used.

BGP Speaker A router that runs the BGP-4 routing protocol. When two BGP speakers are forming an adjacency, they are called “BGP peers” or “neighbors.”

Binding The process of associating a label with a Forwarding Equivalence Class (FEC).

C

CE Device (Customer Edge Device) Router or switch in the customer’s network that is connected to a service provider’s provider edge (PE) router and participates in a Layer 2 or Layer 3 VPNs.

Control Binding Using Control Messages (such as the Label Distribution Protocol) or specific predetermined commands and parameters to bind a label to an FEC. This is a static form of binding.

Customer Edge Device See CE device.

CIDR (Classless Inter-Domain Routing) IP routing which uses the subnet masks to determine the size of each individual subnet. Supports variable subnet lengths. Obsoletes the concept of “Class A, B, C” networks.

Client Peer In BGP route reflection, a member of a cluster that is not the route reflector. See also nonclient peer.

Cluster In BGP, a set of routers that have been grouped together. A cluster consists of one system that acts as a route reflector, along with any number of client peers. The client peers receive their route information only from the route reflector system. Routers in a cluster do not need to be fully meshed.

CR-LDP (Constraint-Based Routing Label Distribution Protocol) Extensions of the Label Distribution Protocol (LDP) to support traffic engineering.

CPE (Customer Premises Equipment) Telephones, routers or other equipment located at a customer site.

Constrained Path In traffic engineering, a path determined using RSVP-TE or CR-LDP signaling and constrained using CSPF. The ERO carried in the packets contains the constrained path information.

CSPF (Constrained Shortest Path First) A Shortest Path First (SPF) IGP algorithm that has been modified to take into account specific restrictions when calculating the shortest path across the network.

D

Data-Driven Binding Dynamically binding a label to an FEC based upon the data stream.

DLCI (Data Link Circuit Identifier) Frame Relay circuit identifiers.

DoD (Downstream on Demand) A method for assigning labels via a label distribution protocol. The ingress router requests a label from the remote end.

DU (Downstream Unsolicited) A method for assigning labels via the Label Distribution Protocol. The egress router assigns labels without being requested to do so from the ingress router.

E

EBGP External BGP Protocol that governs the relationship between BGP speakers belonging to different Autonomous Systems.

EGP (Exterior Gateway Protocol) Any routing protocol used for distributing routes between autonomous systems. BGP-4 is now the most commonly used Exterior Gateway Protocol.

Egress The location at which a datagram exits the MPLS network.

ER (Explicit Route) A route specified at the point of origination. Does not require routing decisions at each hop of the network.

ERO (Explicit Route Object) Extension to RSVP or LDP that allows an RSVP-TE PATH message or CR-LDP Label Request to traverse an explicit sequence of routers that is independent of conventional shortest-path IP routing.

F

Fast Reroute Mechanism for effecting local repair by automatically rerouting traffic from an LSP if a node or link in the LSP fails, thus reducing the loss of packets traveling over the LSP.

FEC (Forwarding Equivalence Class) A group of IP packets forwarded in the same manner – that is, over the same path, with the same priority and the same label.

FIB Forwarding Information Base. A database belonging to a particular router that contains all the necessary information to forward packets from this router to another. The routes installed in the FIB are the best existing routes available at that moment and are used to direct incoming packets to their destination.

I

IBGP Internal BGP. Protocol that governs the relationship between BGP speakers belonging to the same Autonomous System.

IGP (Interior Gateway Protocol) A routing protocol for distributing routes within a single Autonomous System. RIP, OSPF and IS-IS are the most commonly used IGP's.

Ingress The point at which a datagram enters the MPLS network.

IP Prefix A number that uniquely identifies a network or set of networks. The IP prefix of a network consists of the most significant bits that are common to the IP addresses of all nodes belonging to that network.

L

Label A short fixed-length identifier associated with a Forwarding Equivalence Class. (In MPLS, 20-bit, unsigned integer in the range 0 through 1048575 used to identify a packet traveling along an LSP.)

Label Merging The replacement of multiple incoming labels for an FEC with a single outgoing label.

Label Stacking Adding multiple MPLS Labels to a single datagram. This can be used when multiple MPLS networks are traversed and is also used for MPLS VPNs and fast reroute.

Label Swapping Using the incoming label to determine the outgoing label, encapsulation and port. Then replacing the incoming label with the outgoing label.

LER (Label Edge Router) A router at the ingress and/or egress of the MPLS network. This router assigns and/or removes the datagram's label.

LSA Link state advertisement. Unit of information that is advertised by an OSPF router to the other routers in the network. The Link State Advertisement reflects the state of the router interfaces and adjacencies.

LDP (Label Distribution Protocol) A protocol used for distributing MPLS labels within MPLS. There are multiple types of label distribution protocols of which “LDP” is only one possible choice.

LDP Peers Two routers that exchange information with each other via the Label Distribution Protocol.

LSP (Label Switched Path) A data-forwarding path through one or more LSRs determined and based upon the labels attached to each data packet.

LSP Tunnel A traffic engineering LSP capable of carrying multiple data flows.

LSR (Label Switch Router) An MPLS-capable router.

Loose In the context of traffic engineering, a path that can use any route or any number of other intermediate (transit) points to reach the next address in the path. (Definition from RFC 791, modified to fit LSPs.)

M

Martini A method of transporting Layer 2 frames over an MPLS or IP network. Defined in draft-martini-l2circuit-encap-mpls-xx.

MPBGP (Multi-Protocol BGP) An extension to BGP that allows advertisement of IPv6, multicast, VPN-IPv4 and other non-IPv4 topologies within and between BGP autonomous systems.

MPLS (Multi-Protocol Label Switching) A set of protocols for the Internet, which facilitate packet forwarding based on simple labels for connection-oriented capabilities.

MPLS Domain A contiguous set of nodes that operate MPLS routing and forwarding and are found in one Routing or Administrative Domain.

MTU (Maximum Transfer Unit) Limit on segment size for a network.

O

OSPF (Open Shortest Path First) A commonly used interior gateway protocol.

Overlay The process of running MPLS (Label Switching) over an ATM (cell switching) network.

P

Path Attribute Information associated with a particular route. It describes the characteristics of the route and is used to determine the best path to a destination. The Path Attributes are sent as part of BGP protocol messages.

Penultimate LSR The LSR in a traffic flow immediately prior to the egress router. This LSR can also remove the Label from the packet prior to forwarding the datagram to the egress router. This is used to simplify the egress router’s job.

PE Router (Provider Edge Router) A router in the service provider’s network connected to a customer edge (CE) device and that participates in a Virtual Private Network (VPN).

Provider Router (or P Router) Router in the service provider’s network that does not attach to a customer edge (CE) device.

PPVPN (Provider Provisioned Virtual Private Network)

R

RIB Routing Information Base. A database or collection of databases belonging to a particular router that contains the image of the network topology. The information in a router’s RIB is assembled from incoming route advertisements, and it is used to calculate the best paths to all known destinations for installment in the FIB.

RIP (Routing Information Protocol) Another commonly used interior gateway protocol. Typically used for smaller or simpler networks, while OSPF supports large complex infrastructures.

Route Distinguisher Identifier attached to a routing update that distinguishes which VPN it belongs to. Each VPN must have a unique distinguisher associated with it. Each route distinguisher is a 6-byte value.

Route Flapping The repeated advertisement and withdrawal of a route or set of routes due to network instability.

RSVP (Resource Reservation Protocol) A signaling protocol that reserves resources throughout an IP network. Supports IP QoS.

RSVP-TE (Resource Reservation Protocol with Traffic Engineering) One of the two commonly used signaling protocols for the establishment, maintenance and removal of LSPs.

S

S-Bit (Bottom of Stack Bit) A single bit in the MPLS header that indicates the last label in the packet.

Shim Another term for an MPLS label inserted between the layer-2 and layer-3 headers of a packet.

SLA (Service Level Agreement) An agreement (usually a contract) between a service provider and a customer. Guarantees a certain quantitative and/or qualitative level of service.

Static Path In the context of traffic engineering, a static route that requires hop-by-hop manual configuration. No signaling is used to create or maintain the path. Also called a static LSP.

Strict In the context of traffic engineering, a route that must go directly to the next address in the path. (Definition from RFC 791, modified to fit LSPs.)

T

TLV (Type-Length-Value) An encoding method for protocol messages.

Traffic Engineering Process of selecting the paths chosen by data traffic to balance the traffic load on the various links, routers, and switches in the network. (Definition from <http://www.ietf.org/internet-drafts/draft-ietf-mpls-framework-04.txt>.) See also MPLS.

Transit Router In MPLS, any intermediate router in the LSP between the ingress router and the egress router.

Tunnel Private, secure path through an otherwise public network.

V

VLAN Virtual LAN (IEEE 802.1Q) A network whose elements behave as if they are connected to the same physical LAN even though they might be located on separate physical networks.

VoMPLS (Voice over MPLS) A method for carrying voice traffic over an MPLS network.

VPI/VCI (Virtual Path Identifier/Virtual Channel Identifier) ATM channel identifiers.

VPN (Virtual Private Network) A private network created by utilizing shared resources within a public network.

VRF (VPN Routing and Forwarding) The forwarding table contained within a PE router for Layer 3 VPN support.

Spirent Test Methodologies

Spirent Communications has developed a series of test methodology journals to make network and equipment testing easier. Spirent also provides a Web site for scripts and other resources at <http://scdn.spirentcom.com/welcome.asp>. Registration is required to access scripts and other tools to help you expedite testing.

In addition to edge router testing information provided in this edition, the journals offer test methodologies for the following technologies:

Volume I — Router Performance Testing

Volume II — IPv6, Wireless, Edge Router, Metro Optical, VoIP, SS7, QoS

Special IPv6 Edition — IPv6 and IPv4

Optical Edition — Performance Monitoring Statistics

MPLS Edition — MPLS Network Testing



Analyze | Assure | Accelerate™

Spirent Communications is a worldwide provider of integrated performance analysis and service assurance systems for next-generation network technologies. Our solutions accelerate the profitable development and deployment of network equipment and services by emulating real-world conditions in the lab and assuring end-to-end performance of large-scale networks.

Spirent performance analysis solutions include instruments and systems that measure and analyze the performance of network equipment, particularly the devices that route voice and data messages to their destination. Our service assurance solutions include remote test, fault and performance management systems that let network service providers quickly identify network faults and monitor real-time performance.

Spirent's integrated performance analysis and service assurance solutions enable our customers to more rapidly develop and certify new devices, lowering the cost of widespread deployment and operation of new network services.

Spirent Communications is a wholly owned subsidiary of Spirent plc, an international network technology company.

Spirent Communications

15200 Omega Drive
Rockville, MD USA
20850-3240
Tel: +1 301.417.1224
info@spirentcom.com

Sales and Information

Americas
Tel: +1 800.927.2660
productinfo@spirentcom.com

Europe, Middle East, Africa
Tel: +33 1 6137.2250
salesemea@spirentcom.com

Asia Pacific
Tel: +852.2166.8382
spirentasia@spirentcom.com

Spirent Communications Developers Network

<http://scdn.spirentcom.com>

www.spirentcom.com