Network Working Group
Internet Draft
Expires: March 2007

R. Papneja
Isocore
S.Vapiwala
J.Karthik
Cisco Systems
S. Poretsky
Reef Point
S. Rao
Qwest Communications
Jean-Louis Le Roux
France Telecom
October 06

Methodology for benchmarking MPLS Protection mechanisms
<draft-ietf-bmwg-protection-meth-00.txt>

Status of this Memo

Abstract

This draft provides the methodology for benchmarking MPLS Protection mechanisms especially the failover time of local protection (MPLS Fast Reroute as defined in RFC-4090). The failover to a backup tunnel could happen at the headend of the primary tunnel or a midpoint and the backup could offer link or node protection. It becomes vital to benchmark the failover time for all the cases and combinations. The failover time could also greatly differ based on the design and implementation and by factors like the number of prefixes carried by the tunnel, the routing protocols that installed these prefixes (IGP, BGP...), the number of primary tunnels affected by the event that caused the failover, number of primary tunnels the backup protects and type of failure, the physical media type on which the failover occurs etc. All the required benchmarking criteria and benchmarking topology required for measuring failover time of local protection is described Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table of Contents

1. Introduction

A link or a node failure could occur at the headend or the mid point node of a given primary tunnel. The time it takes to failover to the backup tunnel is a key measurement since it directly affects the traffic carried over the tunnel. The failover could occur at the headend or the midpoint of a primary tunnel and the time it takes to failover depends on a variety of factors like the type of physical media, method of FRR solution (detour vs facility), number of primary tunnels, number of

prefixes carried over the tunnel etc. Given all this service providers certainly like to see a methodology to measure the failover time under all possible conditions.

The following sections describe all the different topologies and scenarios that should be used and considered to effectively benchmark the failover time. The failure triggers, procedures, scaling considerations and reporting format of the results are discussed as well.

In order to benchmark failover time, data plane traffic is used as mentioned in [IGP-METH] since traffic loss is measured in a black-box test and is a widely accepted way to measure convergence.

Important point to be noted when benchmarking the failover time is that depending on whether PHP is happening (whether or not implicit null is advertised by the tail-end), and on the number of hops of primary and backup tunnel, we could have different situations where the packets switched over to the backup tunnel may have one, more or 0 labels.

All the benchmarking cases mentioned in this document could apply to facility backup as well as local protection enabled in the detour mode. The test cases and the procedures described here should completely benchmark the failover time of a device under test in all possible scenarios and configuration.

The additional scenarios defined in this document, are in addition to those considered in [FRR-METH]. All the cases enlisted in this document could be verified in a single topology that is similar to this.

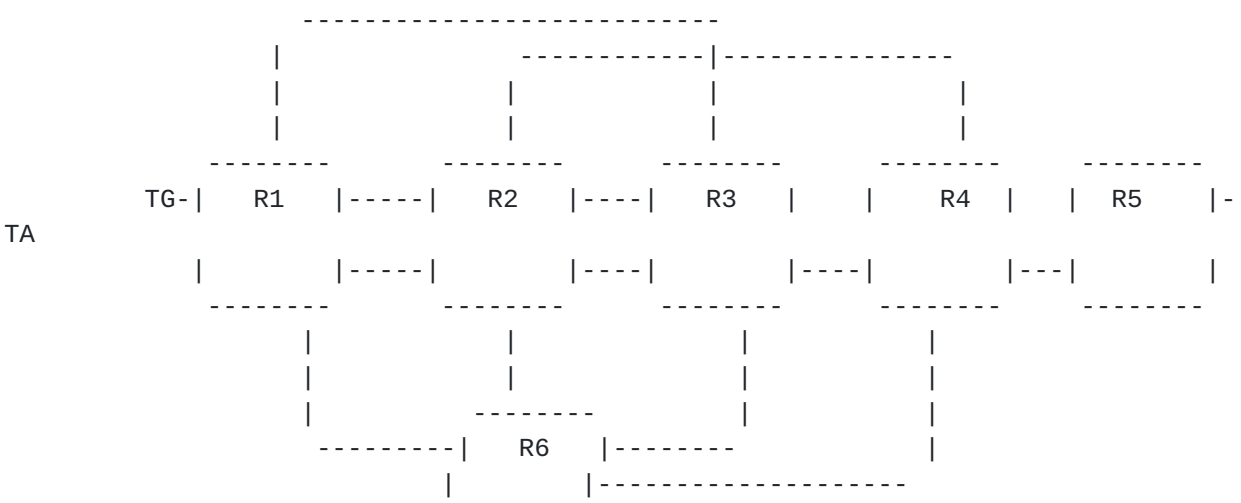

Fig.1: Fast Reroute Topology.

In figure 1, TG & TA are Traffic Generator & Analyzer respectively. A tester is set outside the node as it sends and receives IP traffic along the working Path, run protocol emulations simulating real world peering scenarios. The tester MUST record the number of lost packets, duplicate packet count, reordered packet count, departure time, and arrival time so that the metrics of Failover Time, Additive Latency, and Reversion Time can be measured. The tester may be a single device or a test system.

Two or more failures are considered correlated if those failures occur more or less simultaneously. Correlated failures are often expected where two or more logical resources, such as layer-2 links, rely on a common physical resource, such as common transport. TDM and WDM provide multiplexing at layer-2 and layer-1 that are often the cause of correlated failures. Where such correlations are known, such as knowing that two logical links share a common fiber segment, the expectation of a common failure can be compensated for by specifying Shared Risk Link Groups [RFC-4090]. Not all correlated failures are anticipated in advance of their occurrence. Failures due to natural disasters or due to certain man-made disasters or mistakes are the most notable causes. Failures of this type occur many times a year and generally a quite spectacular failure occurs every few years.

There are two factors impacting service availability. One is the frequency of failure. The other is the duration of failure. FRR improves availability by minimizing the duration of the most common failures. Unexpected correlated failures are less common. Some routers recover much more quickly than others and therefore benchmarking this type of failure may also be useful. Benchmarking of unexpected correlated failures should include measurement of restoration with and without the availability of IP fallback. The use BGP free core may be growing, making the latter case an important test case. This document focuses on FRR failover benchmarking with MPLS TE. Benchmarking of unexpected correlated failures is out of scope but may be covered by a later document.

2. Existing definitions

For the sake of clarity and continuity this RFC adopts the template for definitions set out in Section 2 of RFC 1242. Definitions are indexed and grouped together in sections for ease of reference.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

The reader is assumed to be familiar with the commonly used MPLS terminology, some of which is defined in [MPLS-RSVP], [MPLS-RSVP-TE], and [MPLS-FRR-EXT].

3. Test Considerations

This section discusses the fundamentals of MPLS Protection testing:

- -The types of network events that causes failover
- -Indications for failover
- -the use of data traffic
- -Traffic generation
- -LSP Scaling
- -Reversion of LSP
- -IGP Selection

3.1. Failover Events

Triggers for failover to a backup tunnel are link and node failures seen downstream of the PLR as follows.

Link failure events

- Shutdown interface on PLR side with POS Alarm
- Shutdown interface on remote side with POS Alarm
- Shutdown interface on PLR side with RSVP hello
- Shutdown interface on remote side with RSVP hello
- Shutdown interface on PLR side with BFD
- Shutdown interface on remote side with BFD
- Fiber Pull on PLR side (Both TX & RX or just the Tx)

- Fiber Pull on remote side (Both TX & RX or just the Rx)

- OIR on PLR side
- OIR on remote side
- Sub-interface failure (shutting down of a VLAN)
- Shut parent interface bearing multiple sub-interfaces

Node failure events
A Reload is a graceful shutdown or a power failure. We refer to Crash as a software failure or an assert.

- Reload protected Node, when RSVP Hello are enable
- Crash Protected Node, when RSVP Hello are enable
- Reload Protected Node, when BFD is enable
- Crash Protected Node, when BFD is enable

3.2. Failure Detection [TERMID]

Local failures can be detected via SONET/SDH failure with directly connected LSR. Failure indication may vary with the type of alarm - LOS, AIS, or RDI. Failures on Ethernet technology links such as Gigabit Ethernet rely upon Layer 3 signaling indication for failure.

Different MPLS protection mechanisms and different implementations use different failure indications such as RSVP hellos, BFD etc. Ethernet technologies such as Gigabit Ethernet rely upon layer 3 failure indication mechanisms since there is no Layer 2 failure indication mechanism. The failure detection time may not always be negligible and it could impact the overall failover time.

The test procedures in this document can be used against a local failure as well as against a remote failure to account for completeness of benchmarking and to evaluate failover performance independent of the implemented signaling indication mechanism.

3.3. Use of Data Traffic for MPLS Protection Benchmarking

Customers of service providers use packet loss as the metric for failover time. Packet loss is an externally observable event having direct impact on customers' application performance. MPLS protection

mechanism is expected to minimize the packet loss in the event of a failure. For this reason it is important to develop a standard router benchmarking methodology for measuring MPLS protection that uses

packet loss as a metric. At a known rate for forwarding, packet loss can be measured and used to calculate the Failover time. Measurement of control plane signaling to establish backup paths is not enough to verify failover. Failover is best determined when packets are actually traversing the backup path.

An additional benefit of using packet loss for calculation of Failover time is that it enables black-box tests to be designed. Data traffic can be offered at line-rate to the device under test (DUT), an emulated network event as described above can be forced to occur, and packet loss can be externally measured to calculate the convergence time. Knowledge of DUT architecture is not required. There is no need to rely on the understanding of the implementation details of the DUT to get the required test results.

In addition, this methodology will consider the errored packets and duplicate packets that could have been generated during the failover process. In extreme cases, where measurement of errored and duplicate packets is difficult, these packets could be attributed to lost packets.

3.4. LSP and Route Scaling

Failover time performance may vary with the number of established primary and backup LSPs and routes learned. However the procedure outlined here may be used for any number of LSPs, L, and number of routes protected by PLR, R. L and R must be recorded.

3.5. Selection of IGP

The underlying IGP could be ISIS-TE or OSPF-TE for the methodology proposed here.

3.6. Reversion [TERMID]

Fast Reroute provides a method to return or restore a backup path to original primary LSP upon recovery from the failure. This is referred to as Reversion, which can be implemented as Global Reversion or Local Reversion. In all test cases listed here Reversion should not produce any packet loss, out of order or duplicate packets. Each of the test cases in this methodology document provides a step to verify

that there is no packet loss.

3.7. Traffic generation

It is suggested that there be one or more traffic streams as long as there is a steady and constant rate of flow for all the streams. In order to monitor the DUT performance for recovery times a set of route prefixes should be advertised before traffic is sent. The traffic should be configured towards these routes.

A typical example would be configuring the traffic generator to send the traffic to the first, middle and last of the advertised routes. (First, middle and last could be decided by the numerically smallest, median and the largest respectively of the advertised prefix). Generating traffic to all of the prefixes reachable by the protected tunnel (probably in a Round-Robin fashion, where the traffic is destined to all the prefixes but one prefix at a time in a cyclic manner) is not recommended. The reason why traffic generation is not recommended in a Round-Robin fashion to all the prefixes, one at a time is that if there are many prefixes reachable through the LSP the time interval between 2 packets destined to one prefix may be significantly high and may be comparable with the failover time being measured which does not aid in getting an accurate failover measurement.

3.8. Motivation for topologies

Given that the label stack is dependent on the following 3 entities it is recommended that the benchmarking of failover time be performed on all the 8 topologies enlisted in section 4

- Type of protection (Link Vs Node)

- # of remaining hops of the primary tunnel from the PLR

- # of remaining hops of the backup tunnel from the PLR

4. Test Setup

Topologies to be used for benchmarking the failover time:

This section proposes a set of topologies that covers the scenarios for local protection. All of these 8 topologies shown (figure 2-

figure 9) can be mapped to the master FRR topology shown in figure 1. Topologies shown in section 4.1 to 4.8 refer to the network topologies required to benchmark failover time when DUT is configured as a PLR either in headend or midpoint role. The number of labels listed below are all w.r.t the PLR.

The label stacks shown below each figure in section 4.1 to 4.9 considers the scenario when PHP is enabled.

In the following network topologies,

HE is Head-End, TE is Tail-End, MID is Mid point, MP is Merge Point,

PLR is Point of Local Repair, PRI is Primary and BKP denotes Backup Node

4.1. Link Protection with 1 hop primary (from PLR) and 1 hop backup

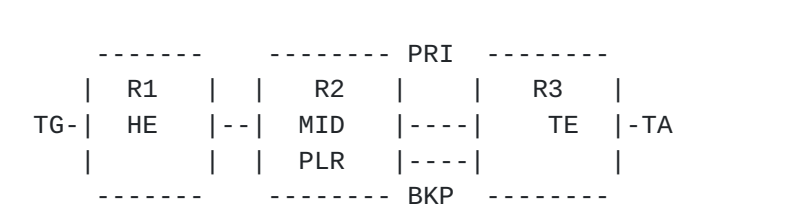

Figure 2: Represents the setup for section 4.1

Traffic	No of Labels before failure	No of labels after failure
IP TRAFFIC (P-P)	0	0
Layer3 VPN (PE-PE)	1	1
Layer3 VPN (PE-P)	2	2
Layer2 VC (PE-PE)	1	1
Layer2 VC (PE-P)	2	2
Mid-point LSPs	0	0

4.2. Link Protection with 1 hop primary (from PLR) and 2 hop backup TE tunnels

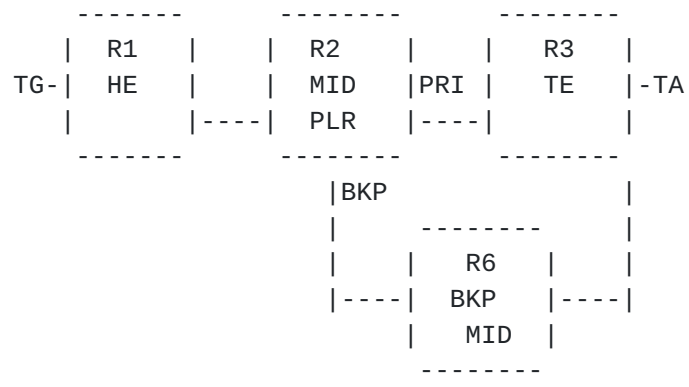

Figure 3: Representing setup for section 4.2

Traffic	No of Labels before failure	No of labels after failure
IP TRAFFIC (P-P)	0	1
Layer3 VPN (PE-PE)	1	2
Layer3 VPN (PE-P)	2	3
Layer2 VC (PE-PE)	1	2
Layer2 VC (PE-P)	2	3
Mid-point LSPs	0	1

4.3. Link Protection with 2+ hop (from PLR) primary and 1 hop backup TE tunnels

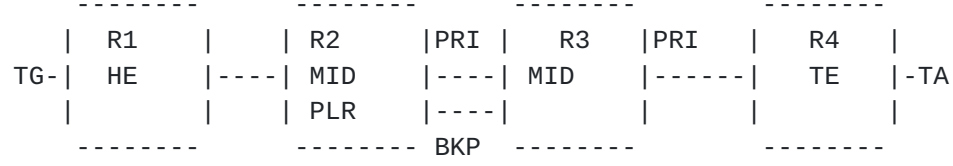

Figure 4: Representing setup for section 4.3

Traffic	No of Labels	No of labels

	before failure	after failure
IP TRAFFIC (P-P)	1	1
Layer3 VPN (PE-PE)	2	2
Layer3 VPN (PE-P)	3	3
Layer2 VC (PE-PE)	2	2
Layer2 VC (PE-P)	3	3
Mid-point LSPs	1	1

4.4. Link Protection with 2+ hop (from PLR) primary and 2 hop backup TE tunnels

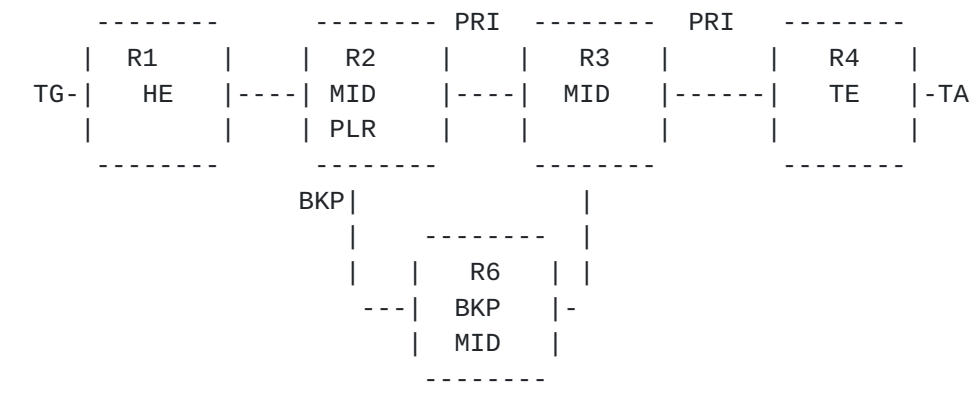

Figure 5: Representing the setup for section 4.4

Traffic	No of Labels before failure	No of labels after failure
IP TRAFFIC (P-P)	1	2
Layer3 VPN (PE-PE)	2	3
Layer3 VPN (PE-P)	3	4
Layer2 VC (PE-PE)	2	3
Layer2 VC (PE-P)	3	4
Mid-point LSPs	1	2

4.5. Node Protection with 2 hop primary (from PLR) and 1 hop backup TE tunnels

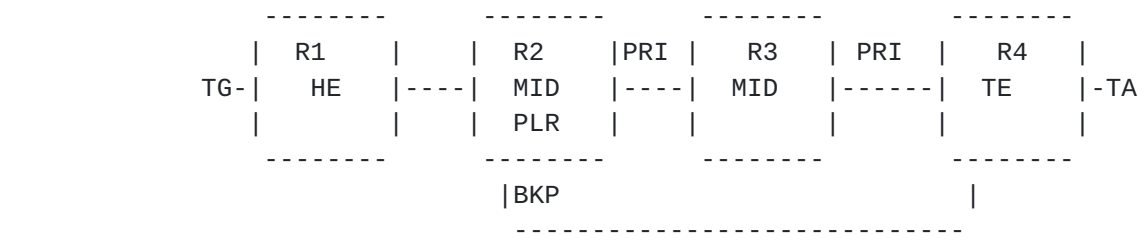

Figure 6: Representing the setup for section 4.5

Traffic	No of Labels before failure	No of labels after failure
IP TRAFFIC (P-P)	1	0
Layer3 VPN (PE-PE)	2	1
Layer3 VPN (PE-P)	3	2
Layer2 VC (PE-PE)	2	1
Layer2 VC (PE-P)	3	2
Mid-point LSPs	1	0

4.6. Node Protection with 2 hop primary (from PLR) and 2 hop backup TE tunnels

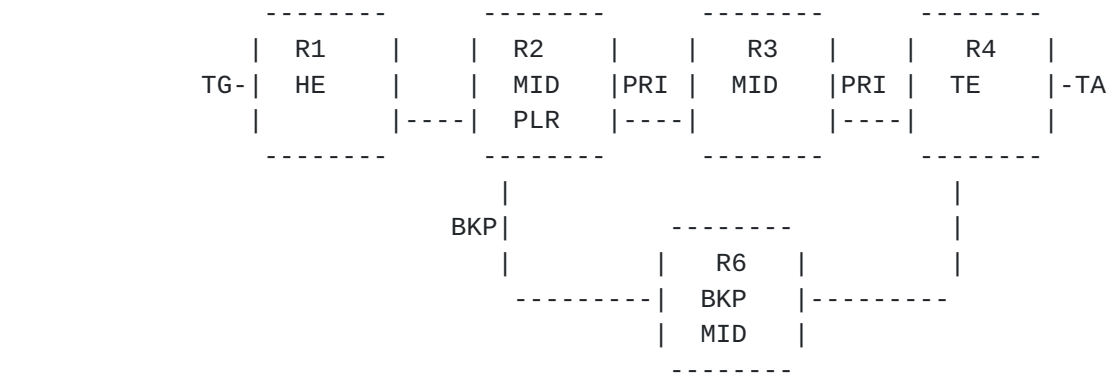

Figure 7: Representing setup for section 4.6

Traffic	No of Labels before failure	No of labels after failure
IP TRAFFIC (P-P)	1	1
Layer3 VPN (PE-PE)	2	2
Layer3 VPN (PE-P)	3	3
Layer2 VC (PE-PE)	2	2
Layer2 VC (PE-P)	3	3
Mid-point LSPs	1	1

4.7. Node Protection with 3+ hop primary (from PLR) and 1 hop backup TE tunnels

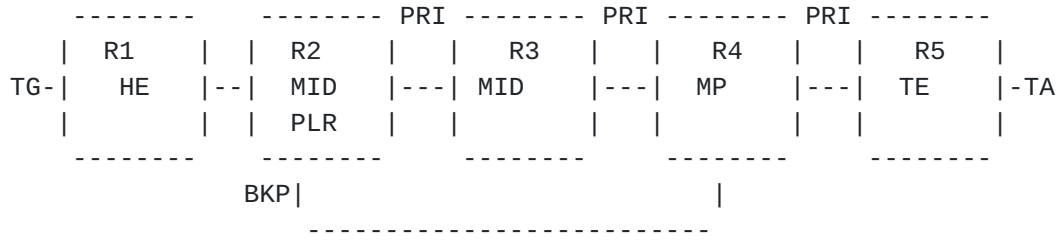

Figure 8: Representing setup for section 4.7

Traffic	No of Labels before failure	No of labels after failure
IP TRAFFIC (P-P)	1	1
Layer3 VPN (PE-PE)	2	2
Layer3 VPN (PE-P)	3	3
Layer2 VC (PE-PE)	2	2
Layer2 VC (PE-P)	3	3
Mid-point LSPs	1	1

4.8. Node Protection with 3+ hop primary (from PLR) and 2 hop backup TE tunnels

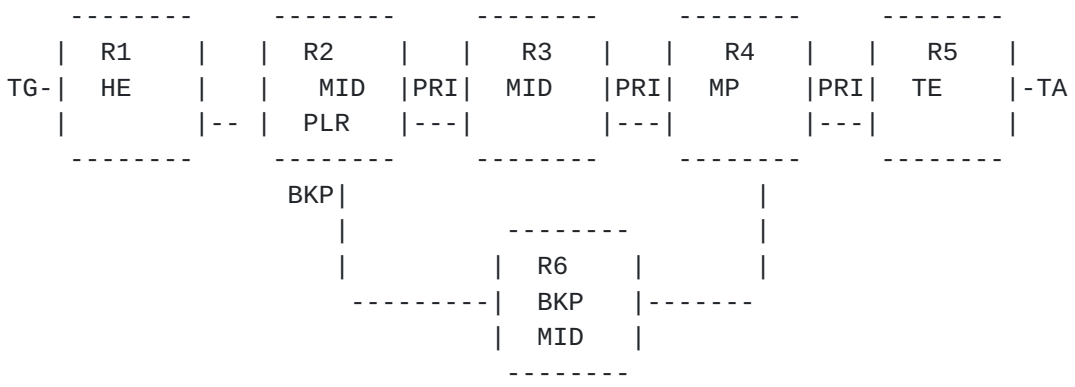

Figure 9: Representing setup for section 4.8

Traffic	No of Labels before failure	No of labels after failure
IP TRAFFIC (P-P)	1	2
Layer3 VPN (PE-PE)	2	3
Layer3 VPN (PE-P)	3	4
Layer2 VC (PE-PE)	2	3
Layer2 VC (PE-P)	3	4
Any	1	2

4.9. Baseline MPLS Forwarding Performance Test Topology

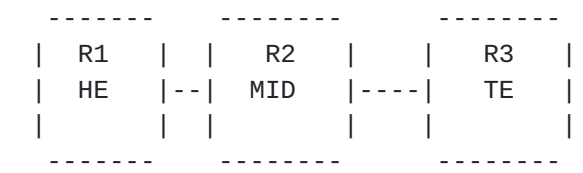

Figure 10: Baseline Forwarding Performance

5. Test Methodology

The procedure described in this section can be applied to all the 8 base test cases and the associated topologies. The backup as well as the primary tunnel are configured to be alike in terms of bandwidth usage. In order to benchmark failover with all possible label stack depth applicable as seen with current deployments, it is suggested that the methodology includes all the scenarios listed here

5.1. Headend as PLR with link failure

Objective

To benchmark the MPLS failover time due to Link failure events described in section 3.1 experienced by the DUT which is the point of local repair (PLR).

Test Setup

- select any one topology out of 8 from section 4
- select overlay technology for FRR test e.g IGP,VPN,or VC
- The DUT will also have 2 interfaces connected to the traffic Generator/analyzer. (If the node downstream of the PLR is not A simulated node, then the Ingress of the tunnel should have one link connected to the traffic generator and the node downstream to the PLR or the egress of the tunnel should have a link connected to the traffic analyzer).

Test Configuration

1. Configure the number of primaries on R2 and the backups on R2 as required by the topology selected.
2. Advertise prefixes (as per FRR Scalability table describe in Appendix A) by the tail end.

Procedure

1. Establish the primary lsp on R2 required by the topology selected

2. Establish the backup lsp on R2 required by the selected topology

3. Verify primary and backup lsps are up and that primary is protected
4. Verify Fast Reroute protection is enabled and ready
5. Setup traffic streams as described in section 3.7
6. Send IP traffic at maximum Forwarding Rate to DUT.
7. Verify traffic switched over Primary LSP.
8. Trigger any choice of Link failure as describe in section 3.1
9. Verify that primary tunnel and prefixes gets mapped to backup tunnels
10. Stop traffic stream and measure the traffic loss.
11. Failover time is calculated as per defined in section 6, Reporting format.
12. Start traffic stream again to verify reversion when protected interface comes up. Traffic loss should be 0 due to make before break or reversion.
13. Enable protected interface that was down (Node in the case of NNHOP)
14. Verify head-end signals new LSP and protection should be in place again

5.2. Mid-Point as PLR with link failure

Objective

To benchmark the MPLS failover time due to Link failure events described in section 3.1 experienced by the device under test which is the point of local repair (PLR).

Test Setup

- select any one topology out of 8 from section 4
- select overlay technology for FRR test as Mid-Point lsps
- The DUT will also have 2 interfaces connected to the traffic generator.

Test Configuration

1. Configure the number of primaries on R1 and the backups on R2 as required by the topology selected

2. Advertise prefixes (as per FRR Scalability table describe in Appendix A) by the tail end.

Procedure

1. Establish the primary lsp on R1 required by the topology selected
2. Establish the backup lsp on R2 required by the selected topology
3. Verify primary and backup lsps are up and that primary is protected
4. Verify Fast Reroute protection
5. Setup traffic streams as described in section 3.7
6. Send IP traffic at maximum Forwarding Rate to DUT.
7. Verify traffic switched over Primary LSP.
8. Trigger any choice of Link failure as describe in section 3.1
9. Verify that primary tunnel and prefixes gets mapped to backup tunnels
10. Stop traffic stream and measure the traffic loss.
11. Failover time is calculated as per defined in section 6, Reporting format.
12. Start traffic stream again to verify reversion when protected interface comes up. Traffic loss should be 0 due to make before break or reversion
13. Enable protected interface that was down (Node in the case of NNHOP)
14. Verify head-end signals new LSP and protection should be in place again

5.3. Headend as PLR with Node failure

Objective

To benchmark the MPLS failover time due to Node failure events described in section 3.1 experienced by the device under test which is the point of local repair (PLR).

Test Setup

- select any one topology from section 4.5 to 4.8
- select overlay technology for FRR test e.g IGP,VPN,or VC
- The DUT will also have 2 interfaces connected to the traffic generator.

Test Configuration

1. Configure the number of primaries on R2 and the backups on R2 as required by the topology selected
2. Advertise prefixes (as per FRR Scalability table describe in Appendix A) by the tail end.

Procedure

1. Establish the primary lsp on R2 required by the topology selected
2. Establish the backup lsp on R2 required by the selected topology
3. Verify primary and backup lsps are up and that primary is protected
4. Verify Fast Reroute protection
5. Setup traffic streams as described in section 3.7
6. Send IP traffic at maximum Forwarding Rate to DUT.
7. Verify traffic switched over Primary LSP.
8. Trigger any choice of Node failure as describe in section 3.1
9. Verify that primary tunnel and prefixes gets mapped to backup tunnels
10. Stop traffic stream and measure the traffic loss.
11. Failover time is calculated as per defined in section 6, Reporting format.
12. Start traffic stream again to verify reversion when protected interface comes up. Traffic loss should be 0 due to make before break or reversion
13. Boot protected Node that was down.
14. Verify head-end signals new LSP and protection should be in place again

5.4. Mid-Point as PLR with Node failure

Objective

To benchmark the MPLS failover time due to Node failure events described in section 3.1 experienced by the device under test which is the point of local repair (PLR).

Test Setup

- select any one topology from section 4.5 to 4.8
- select overlay technology for FRR test as Mid-Point lsps
- The DUT will also have 2 interfaces connected to the traffic generator.

Test Configuration

1. Configure the number of primaries on R1 and the backups on R2 as required by the topology selected
2. Advertise prefixes (as per FRR Scalability table describe in Appendix A) by the tail end.

Procedure

1. Establish the primary lsp on R1 required by the topology selected
2. Establish the backup lsp on R2 required by the selected topology
3. Verify primary and backup lsps are up and that primary is protected
4. Verify Fast Reroute protection
5. Setup traffic streams as described in section 3.7
6. Send IP traffic at maximum Forwarding Rate to DUT.
7. Verify traffic switched over Primary LSP.
8. Trigger any choice of Node failure as describe in section 3.1
9. Verify that primary tunnel and prefixes gets mapped to backup tunnels
10. Stop traffic stream and measure the traffic loss.
11. Failover time is calculated as per defined in section 6, Reporting format.

12. Start traffic stream again to verify reversion when protected interface comes up. Traffic loss should be 0 due to make before break or reversion
13. Boot protected Node that was down
14. Verify head-end signals new LSP and protection should be in place again

5.5. Baseline MPLS Forwarding Performance Test Cases

For the following Forwarding Performance Benchmarking cases, the egress must not send an implicit-null label. That is PHP should not occur.

5.5.1. DUT Throughput as Ingress

Objective

To baseline the MPLS Throughput of the DUT acting as an Ingress.

Procedure

1. Configure the DUT as R1, Ingress and the Tester as R2/R3 Midpoint and Egress as shown in Figure 10.
2. Execute the Throughput benchmarking test, as specified in [RFC-BENCH], paragraph 26.1.

Expected Results:

The DUT will push a single label onto the IP packet and forward it to the Tester as an MPLS packet.

5.5.2. DUT Latency as Ingress

Objective

To baseline the MPLS Latency of the DUT acting as an Ingress.

Procedure

1. Configure the DUT as R1, Ingress and the Tester as R2/R3 Midpoint and Egress as shown in Figure 10.
2. Execute the Latency benchmarking test, as specified in [RFC-BENCH], paragraph 26.2.

Expected Results:

The DUT will push a single label onto the IP packet and forward it to the Tester as an MPLS packet.

5.5.3. DUT Throughput as Egress

Objective

To baseline the MPLS Throughput of the DUT acting as an Egress.

Procedure

1. Configure the DUT as R3, Egress and the Tester as R1/R2 Ingress and Midpoint as shown in Figure 10.
2. Execute the Throughput benchmarking test, as specified in [RFC-BENCH], paragraph 26.1 using MPLS labeled IP packets for the offered load.

Expected Results:

The DUT will pop a single label from the IP packet and forward it to the Tester as an IP packet.

5.5.4. DUT Latency as Egress

Objective

To baseline the MPLS Latency of the DUT acting as an Egress.

Procedure

1. Configure the DUT as R3, Egress and the Tester as R1/R2 Ingress and Midpoint as shown in Figure 10.

2. Execute the Latency benchmarking test, as specified in [RFC-BENCH], paragraph 26.2 using MPLS labeled IP packets for the offered load.

Expected Results:

The DUT will pop a single label from the IP packet and forward it to the Tester as an IP packet.

5.5.5. DUT Throughput as Mid-Point

Objective

To baseline the MPLS Throughput of the DUT acting as a Mid-Point.

Procedure

1. Configure the DUT as R2, Mid-Point and the Tester as R1/R3 Ingress and Egress as shown in Figure 10.
2. Execute the Throughput benchmarking test, as specified in [RFC-BENCH], paragraph 26.1 using MPLS labeled IP packets for the offered load.

Expected Results:

The DUT will receive the MPLS labeled packet, swap a single MPLS label and forward it to the Tester as an MPLS labeled packet.

5.5.6. DUT Latency as Mid-Point

Objective

To baseline the MPLS Latency of the DUT acting as a Mid-Point.

Procedure

1. Configure the DUT as R2, Mid-Point and the Tester as R1/R3 Ingress and Egress as shown in Figure 10.

2. Execute the Latency benchmarking test, as specified in [RFC-BENCH], paragraph 26.2 using MPLS labeled IP packets for the offered load.

Expected Results:

The DUT will receive the MPLS labeled packet, swap a single MPLS label and forward it to the Tester as an MPLS labeled packet.

6. Reporting Format

For each test, it is recommended that the results be reported in the following format.

Parameter	Units
IGP used for the test	ISIS-TE/ OSPF-TE
Interface types	Gige,POS,ATM,VLAN etc.
Packet Sizes offered to the DUT	Bytes
IGP routes advertised	number of IGP routes
RSVP hello timers configured (if any)	milliseconds
Number of FRR tunnels configured	number of tunnels
Number of VPN routes in head-end	number of VPN routes
Number of VC tunnels	number of VC tunnels
Number of BGP routes	number of BGP routes
Number of mid-point tunnels	number of tunnels
Number of Prefixes protected by Primary	number of prefixes
Number of LSPs being protected	number of LSPs
Topology being used	Section number
Failure Event	Event type

Benchmarks

Minimum failover time	milliseconds
Mean failover time	milliseconds

Maximum failover time	milliseconds
Minimum reversion time	milliseconds
Mean reversion time	milliseconds
Maximum reversion time	milliseconds

Failover time suggested above is calculated using the following formula: (Numbers of packet drop/rate per second * 1000) milliseconds

Note: If the primary is configured to be dynamic, and if the primary is to reroute, make before break should occur from the backup that is in use to a new alternate primary. If there is any packet loss seen, it should be added to failover time.

7. Security Considerations

Documents of this type do not directly affect the security of the Internet or of corporate networks as long as benchmarking is not performed on devices or systems connected to operating networks.

8. Acknowledgements

We would like to thank Jean Philip Vasseur for his invaluable input to the document and Curtis Villamizar for his contribution in suggesting
text on definition and need for benchmarking Correlated failures.

Additionally we would like to thank Arun Gandhi, Amrit Hanspal, Karu Ratnam and for their input to the document.

9. References

9.1. Normative References

[MPLS-RSVP] R. Braden, Ed., et al, "Resource ReSerVation protocol (RSVP) -- version 1 functional specification," RFC2205, September 1999.

[MPLS-RSVP-TE] D. Awduche, et al, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC3209, December 2001.

"Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090.

[MPLS-ARCH] Rosen, E., Viswanathan, A. and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC-BENCH] Bradner, S. and McQuaid, J., "Benchmarking Methodology for Network Interconnect Devices", RFC 2544.

9.2. Informative References

[MPLS-LDP] Andersson, L., Doolan, P., Feldman, N., Fredette, A. and B. Thomas, "LDP Specification", RFC 3036, January 2001.

[RFC-WORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.

[RFC-IANA] T. Narten and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 2434.

[TERM-ID] Poretsky, S., Papneja, R., "Benchmarking Terminology for Protection Performance", draft-poretsky-protection-term-00.txt, work in progress.

[IGP-METH] S. Poretsky, B. Imhoff. "Benchmarking Methodology for IGP Data Plane Route Convergence," draft-ietf-bmwg-igp-dataplane-conv-meth-11.txt, work in progress.

10. Author's Address

Rajiv Papneja

Isocore
12359 Sunrise Valley Drive, STE 100
Reston, VA 20190
USA
Phone: +1 703 860 9273
Email: rpapneja@isocore.com

Samir Vapiwala
Cisco System
300 Beaver Brook Road
Boxborough, MA 01719
USA
Phone: +1 978 936 1484
Email: svapiwal@cisco.com

Jay Karthik
Cisco System
300 Beaver Brook Road
Boxborough, MA 01719
USA
Phone: +1 978 936 0533
Email: jkarthik@cisco.com

Scott Poretsky
Reef Point Systems
8 New England Executive Park
Burlington, MA 01803
USA
Phone: + 1 781 395 5090
EMail: sporetsky@reefpoint.com

Shankar Rao
Qwest Communications,
950 17th Street
Suite 1900
Qwest Communications
Denver, CO 80210
USA
Phone: + 1 303 437 6643
Email: shankar.rao@qwest.com

Jean-Louis Le Roux
France Telecom
2 av Pierre Marzin
22300 Lannion
France
Phone: 00 33 2 96 05 30 20
Email: jeanlouis.leroux@orange-ft.com

Full Copyright Statement

http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement
Funding for the RFC Editor function is currently provided by the Internet Society.

Appendix A: Fast Reroute Scalability Table

This section provides the recommended numbers for evaluating the scalability of fast reroute implementations. It also recommends the typical numbers for IGP/VPNv4 Prefixes, LSP Tunnels and VC entries. Based on the features supported by the device under test, appropriate scaling limits can be used for the test bed.

A 1. FRR IGP Table

No of Headend TE LSPs	IGP Prefixes
1	100
1	500
1	1000
1	2000
1	5000
2(Load Balance)	100
2(Load Balance)	500
2(Load Balance)	1000
2(Load Balance)	2000
2(Load Balance)	5000
100	100
500	500
1000	1000
2000	2000

A 2. FRR VPN Table

No of Headend TE LSPs	VPNv4 Prefixes
1	100
1	500
1	1000
1	2000
1	5000
1	10000
1	20000
1	Max
2(Load Balance)	100
2(Load Balance)	500
2(Load Balance)	1000
2(Load Balance)	2000
2(Load Balance)	5000
2(Load Balance)	10000
2(Load Balance)	20000
2(Load Balance)	Max

A 3. FRR Mid-Point LSP Table

No of Mid-point TE LSps could be configured at the following recommended levels

100
500
1000
2000
Max supported number

A 4. FRR VC Table

No of Headend TE LSPs	VC entries
1	100
1	500
1	1000

1	2000
1	Max
100	100
500	500
1000	1000
2000	2000

Appendix B: Abbreviations

BFD	- Bidirectional Fault Detection
BGP	- Border Gateway protocol
CE	- Customer Edge
DUT	- Device Under Test
FRR	- Fast Reroute
IGP	- Interior Gateway Protocol
IP	- Internet Protocol
LSP	- Label Switched Path
MP	- Merge Point
MPLS	- Multi Protocol Label Switching
N-Nhop	- Next - Next Hop
Nhop	- Next Hop
OIR	- Online Insertion and Removal
P	- Provider
PE	- Provider Edge
PHP	- Penultimate Hop Popping
PLR	- Point of Local Repair
RSVP	- Resource reSerVation Protocol
SRLG	- Shared Risk Link Group
TA	- Traffic Analyzer
TE	- Traffic Engineering
TG	- Traffic Generator
VC	- Virtual Circuit
VPN	- Virtual Private Network