Benchmarking Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 4, 2014

B. Parise
Cisco Systems
R. Papneja
Huawei Technologies
July 3, 2013

Terminology for Benchmarking LDP Data Plane Convergence
draft-parise-bmwg-ldp-convergence-term-00.txt

Abstract

This document defines new terms for benchmarking of LDP convergence. These terms are to be used in future methodology documents for benchmarking LDP Convergence. Existing BMWG terminology documents such as IGP Convergence Benchmarking [RFC 6412] provide useful terms for LDP Convergence benchmarking. These terms are discussed in this document. Applicable terminology for MPLS and LDP defined in MPLS WG RFCs [RFC 3031] and [RFC 5036] are also discussed.

Table of Contents

1. Introduction

This draft describes the terminology for benchmarking LDP Convergence. An accompanying document will describe the methodology for doing the benchmarking. The main motivation for doing this work is the increased focus on lowering convergence time for LDP as an alternative to other solutions such as MPLS Fast Reroute (i.e. protection techniques using RSVP-TE extensions).

The purpose of this document is to find existing terminology as well as define new terminology when needed terms are not available. The terminology will support the methodology that will be based on black-box testing of the LDP dataplane. The approach is very similar to the one found in [RFC 6412] and [RFC 6413].

2. Existing Definitions

2.1. BMWG Convergence Terms

This document uses existing terminology defined in other IETF documents. These include the following:

Route Convergence	Defined in [RFC 6412]
Convergence Packet Loss	Defined in [RFC 6412]
Convergence Event Instant	Defined in [RFC 6412]
Convergence Recovery Instant	Defined in [RFC 6412]
Rate-Derived Convergence Time	Defined in [RFC 6412]
Convergence Event Transition	Defined in [RFC 6412]
Convergence Recovery Transition	Defined in [RFC 6412]
Loss-Derived Convergence Time	Defined in [RFC 6412]
Restoration Convergence Time	Defined in [RFC 6412]
Packet Sampling Interval	Defined in [RFC 6412]
Local Interface	Defined in [RFC 6412]
Neighbor Interface	Defined in [RFC 6412]
Remote Interface	Defined in [RFC 6412]
Preferred Egress Interface	Defined in [RFC 6412]
Next-Best Egress Interface	Defined in [RFC 6412]
Stale Forwarding	Defined in [RFC 6412]

2.2. MPLS/LDP Terms

Label	Defined in [RFC 3031]
FEC	Defined in [RFC 3031]
Label Withdraw	Defined in [RFC 5036]
LSP	Defined in [RFC 3031]
LSR	Defined in [RFC 3031]
LDP Identifier	Defined in [RFC 5036]
LDP Session	Defined in [RFC 5036]
Per-Interface Label Space	Defined in [RFC 3031]
Per-Platform Label Space	Defined in [RFC 3031]
MPLS Node	Defined in [RFC 3031]
MPLS Edge Node	Defined in [RFC 3031]
MPLS Egress Node	Defined in [RFC 3031]
MPLS Ingress Node	Defined in [RFC 3031]
Upstream LSR	Defined in [RFC 3031]
Downstream LSR	Defined in [RFC 3031]
Local Repair	Defined in [RFC 4090]
PLR	Defined in [RFC 4090]
One-to-One Backup	Defined in [RFC 4090]
Detour LSP	Defined in [RFC 4090]
Backup Path	Defined in [RFC 4090]
Downstream-on-Demand	Defined in [RFC 3031]
Unsolicited Downstream	Defined in [RFC 3031]
Independent Label Distribution Control	Defined in [RFC 5036]
Address Family	Defined in [RFC 5036]
IGP Update Message	ISIS/OSPF LSA

3. Term Definitions

3.1. LDP Binding Table

Definition:

Table in which the LSR maintains all learned labels. It consists of the prefix and label information bound to a peer's LDP identifier and the list of sent and received bindings/peer.

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

FEC Forwarding Table

3.2. FEC Forwarding Table

Definition:

Table in which the LSR maintains the next hop information for the particular FEC with the associated outgoing label and interface. The information used for setting up the FEC forwarding table is retrieved from the LDP Binding Table.

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

LDP Binding Table

3.3. FEC Convergence Event

Definition:

The occurrence of a planned or unplanned action in the network that results in a change to an LSR’s LDP next-hop forwarding.

Discussion:

Convergence Events include link loss, routing protocol session loss, router failure, and better next-hop. Also, different types of administrative events such as interface shutdown is considered.

Measurement Units:

N/A

Issues:

None

See Also:

FEC Forwarding Table Convergence

FEC Convergence

3.4. FEC Forwarding Table Convergence

Definition:

Recovery from a FEC Convergence Event that causes the FEC Forwarding Table to change and re-stabilize.

Discussion:

FEC Forwarding Table Convergence updates after the RIB and LDP Binding Table update due to a FEC Convergence Event. FEC Forwarding Table Convergence can be observed externally by the rerouting of data Traffic to a new egress interface.

Measurement Units:

seconds

Issues:

None

See Also:

FEC Forwarding Table

FEC Convergence Event

FEC Convergence

3.5. FEC Convergence

Definition:

Recovery from a FEC Convergence Event that causes the LDP Binding Table to change and re-stabilize.

Discussion:

FEC Convergence is a change in an LDP Binding of a prefix and label to a peer's LDP Identifier. This change can be an update or recovery due to a FEC Convergence Event. FEC Convergence is an LSR action made prior to FEC Forwarding Table Convergence. FEC Convergence is not an externally observable Black-Box measurement.

Measurement Units:

N/A

Issues:

Where is LDP Identifier defined? Where is LDP Binding defined?

See Also:

FEC Binding Table

FEC Convergence Event

FEC Forwarding Table Convergence

3.6. Multiple Next-Hop FEC

Definition:

A FEC with more than one next-hop and associated outgoing label and interface.

Discussion:

A Multiple Next-Hop FEC can be verified from the FEC Forwarding Table and from externally observing traffic being forwarded to a FEC on one or more interfaces.

Measurement Units:

N/A

Issues:

None

See Also:

FEC Forwarding Table

3.7. Ingress LSR

Definition:

An MPLS ingress node which is capable of forwarding native L3 packets.

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

MPLS Node

MPLS Edge Node

MPLS Egress Node

MPLS Ingress Node

Label Switching Router (LSR)

Egress LSR

3.8. Egress LSR

Definition:

An MPLS Egress node which is capable of forwarding native L3 packets.

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

MPLS Node

MPLS Edge Node

MPLS Egress Node

MPLS Ingress Node

Label Switching Router (LSR)

Ingress LSR

3.9. LDP Peer

Definition:

An adjacent LSR with which LDP adjacency is established

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

Targeted LDP Peer

3.10. Targeted LDP Peer

Definition:

An adjacent LSR (usually more than a hop away) with which LDP adjacency is established through a directed hello message which is unicast.

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

LDP Peer

3.11. Targeted FECs

Definition:

The FECs advertised by a Targeted LDP Peer

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

Targeted Peer

3.12. Multi-Labeled Packets

Definition:

A data packet that has more than one label in the label stack.

Discussion:

This typically happens when a Targeted Peer is established over a traffic engineered tunnel.

Measurement Units:

N/A

Issues:

None

See Also:

None

3.13. Equal Cost Multiple Paths

Definition:

Existence of multiple IGP paths to reach a particular destination. In this case the depending on the implementation traffic destined to a prefix that has multiple equal cost paths is load balanced across all these paths.

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

Equal Cost Multiple FECs

3.14. Equal Cost Multiple FECs

Definition:

Existence of multiple to reach a destination. Typically the LSR that has multiple FECs of equal costs does a load balance on all the FECs

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

Equal Cost Multiple Paths

3.15. FEC Convergence at Ingress LSR

Definition:

Recovery from a FEC Convergence Event that causes the LDP Binding Table to change and re-stabilize at the Ingress LSR

Discussion:

FEC Convergence is a change in an LDP Binding of a prefix and label to a peer’s LDP Identifier. This change can be an update or recovery due to a FEC Convergence Event. FEC Convergence is an LSR action made prior to FEC Forwarding Table Convergence. FEC Convergence is not an externally observable Black-Box measurement.

Measurement Units:

N/A

Issues:

Where is LDP Identifier defined? Where is LDP Binding defined?

See Also:

LDP Binding Table

FEC Convergence Event

FEC Forwarding Table Convergence

3.16. FEC Convergence at Midpoint LSR

Definition:

Recovery from a FEC Convergence Event that causes the LDP Binding Table to change and re-stabilize at a Midpoint LSR

Discussion:

FEC Convergence is a change in an LDP Binding of a prefix and label to a peer’s LDP Identifier. This change can be an update or recovery due to a FEC Convergence Event. FEC Convergence is an LSR action made prior to FEC Forwarding Table Convergence. FEC Convergence is not an externally observable Black-Box measurement.

Measurement Units:

N/A

Issues:

Where is LDP Identifier defined? Where is LDP Binding defined?

See Also:

LDP Binding Table

FEC Convergence Event

FEC Forwarding Table Convergence

3.17. LDP Advertisement Type

Definition:

The type of LDP advertisement in operation. Downstream On Demand vs Downstream Unsolicited.

Discussion:

None

Measurement Units:

N/A

Issues:

None

See Also:

None

3.18. Label Merging LSR

Definition:

A LSR which is capable of sending multiple packets out of the same outgoing interface with the same label even though it receives these packets from different incoming interfaces and may also receive them with the same lane

Discussion:

With label merging the LSR need to send a single label per FEC and also on the receiving end the number of incoming labels per FEC is never larger than the number of label distribution adjacencies

Measurement Units:

N/A

Issues:

There maybe be scenarios where a Merging LSR is capable of merging only a subset of incoming labels into a single outgoing label

See Also:

Non-Merging LSR and [RFC 3031]

3.19. Non-merging LSR

Definition:

A LSR which forwards packets with multiple outgoing labels when it receives packets from the same FEC with different incoming labels

Discussion:

Without label merging the number of outgoing labels per FEC could be as large as the number of nodes in the network

Measurement Units:

N/A

Issues:

None

See Also:

Label Merging LSR and [RFC 3031]

3.20. LDPv6

Definition:

This term implies forwarding of IPv6 packets as detailed in [RFC 5036]

Discussion:

None

Measurement Units:

N/A

Issues:

The current specification [RFC 5036] has certain gaps as detailed in [LDPv6]. Once its standardized we will extend the scope to cover those details.

See Also:

None

4. Factors impacting Convergence

4.1. Interaction with Other Protocols

LDP convergence must include the affect of interaction with IGPs. All test reports must include the IGPs provisioned in the test and their associated parameters

4.2. Timers

LDP convergence is impacted by the Hold and Keepalive Timers. Test reports must include all the relevant timer values

4.3. TCP Parameters

As LDP uses TCP for sessions, all relevant TCP session parameters must be reported

5. Security Considerations

Benchmarking activities as described in this memo are limited to technology characterization using controlled stimuli in a laboratory environment, with dedicated address space and the constraints specified in the sections above.

The benchmarking network topology will be an independent test setup and MUST NOT be connected to devices that may forward the test traffic into a production network, or misroute traffic to the test management network.

Further, benchmarking is performed on a "black-box" basis, relying solely on measurements observable external to the DUT/SUT.

Special capabilities SHOULD NOT exist in the DUT/SUT specifically for benchmarking purposes. Any implications for network security arising from the DUT/SUT SHOULD be identical in the lab and in production networks.

6. Acknowledgements

We thank Al Morton for providing valuable comments to this document. We also thank Scott Poretsky for his contributions to the initial version of this document.

7. References

7.1. Normative References

[I-D.ietf-mpls-ldp-ipv6]
Asati, R., Manral, V., Papneja, R., and C. Pignataro, "Updates to LDP for IPv6", draft-ietf-mpls-ldp-ipv6-08 (work in progress), February 2013.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.

[RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", RFC 5036, October 2007.

[RFC6412] Poretsky, S., Imhoff, B., and K. Michielsen, "Terminology for Benchmarking Link-State IGP Data-Plane Route Convergence", RFC 6412, November 2011.

[RFC6413] Poretsky, S., Imhoff, B., and K. Michielsen, "Benchmarking Methodology for Link-State IGP Data-Plane Route Convergence", RFC 6413, November 2011.

7.2. Informative References

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.

Authors' Addresses

Bhavani Parise
Cisco Systems

Email: bhavani@cisco.com

Rajiv Papneja
Huawei Technologies

Email: rajiv.papneja@huawei.com