

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 1, 2021

J. He, Ed.
B. Nie
Z. Zhao
Ericsson
June 30, 2020

MPLS Entropy Block
draft-jiang-mpls-entropy-block-01

Abstract

Load balancing across MPLS networks requires entropy information to be carried in MPLS label stack and to be handled by the downstream Label Switching Routers. This document describes the problems using existing approaches and introduces an approach to the solution.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 1, 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Internet-Draft

MPLS Entropy Block

June 2020

Table of Contents

1.	Introduction	2
2.	Terminology	3
3.	Approach	3
4.	End-to-End Processing	3
4.1.	Example 1 Entropy Block advertisement through signaling .	3
4.2.	Example 2 Entropy Block advertisement through global configuration	4
5.	Signaling for Entropy Block	5
6.	Issues	5
6.1.	Anycast Segment	5
7.	IANA Considerations	6
8.	Security Considerations	6
9.	References	6
9.1.	Normative References	6
9.2.	Informative References	7
	Authors' Addresses	7

[1.](#) Introduction

Traffic load balancing improves bandwidth utilization in MPLS networks. The entropy information is required to be carried in MPLS label stack of each packet in the MPLS network, and to be handled along the path by Label Switching Routers (LSR) towards egress node. [\[RFC6790\]](#) provides a technique used in the MPLS data plane to provide entropy for load-balancing, using a combination of Entropy Label Indicator and Entropy Label. [\[RFC6391\]](#) provides another technique used in the MPLS data plane while only serve for pseudowire based services.

An LSR may have limitations in the number of labels it can push. It may also have a limitation in the number of labels it can inspect when looking for entropy information. The limitations become obvious for traffic-engineering applications such as Segment Routing [\[RFC8402\]](#) and Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) [\[RFC3209\]](#). Some proposals are discussed in [\[RFC8662\]](#).

This document describes an approach to mitigate the limitations, by allocating a specific MPLS label block for entropy usage. Thus the entropy information can be carried in the MPLS label stack and handled by the downstream LSRs. It adds additional alternatives beyond those discussed in [\[RFC6790\] section 2](#). This single entropy

label approach improves efficiency when compared with [[RFC6790](#)] which uses label pair (ELI/EL). It also accommodates label location flexibility and service agnostic when compared with [[RFC6391](#)]. Considerations from [[RFC8662](#)] are still applicable, whether it is ELI/EL or EL only. As described there, systems involved still need

to carefully consider issues such as MSD, and determine how many entropy labels are needed, and where they need to be placed. This draft reduces the overhead of these labels from 2 labels to 1.

In the document we see some issues with all the alternatives and are looking for further discussion.

[2.](#) Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

[3.](#) Approach

The entropy information is encoded as an additional label in MPLS label stack. Since the single Entropy Label (EL) is encapsulated by an ingress LSR while terminated at an egress LSR, both ingress LSR and egress LSR MUST be able to distinguish unambiguously between entropy label and other labels. To accomplish this, a dedicated label range is allocated for entropy label, named Entropy Block (EB).

The Entropy Block can be advertised through signals (E.g. through IGP) or globally configured (E.g. through NMS).

On ingress LSR, when encoding entropy information in MPLS label stack towards an egress LSR, the label value shall be in the range of the Entropy Block. When received on the egress LSR and the entropy label appears from the MPLS label stack, it is simply popped.

[4.](#) End-to-End Processing

[4.1.](#) Example 1 Entropy Block advertisement through signaling

In this example, Entropy Block is advertised by each LSR through IGP (E.g. IS-IS). Each LSR may have different Entropy Block. One traffic flow is forwarded through a Traffic Engineering (TE) tunnel, LSR1->LSR4->LSR6.

On ingress LSR, entropy value is calculated from packet received, and adjusted to the Entropy Block of LSR6 as it is to be popped by LSR6.

On transit LSRs, such as LSR2 or LSR3, the entropy information in the packet is utilized for load balancing. When packet reaches LSR4, the steering point of the TE tunnel, the topmost tunnel label is popped

and packet continues to egress LSR, LSR6. After the second tunnel label is popped, the entropy label is popped and packet is sent out.

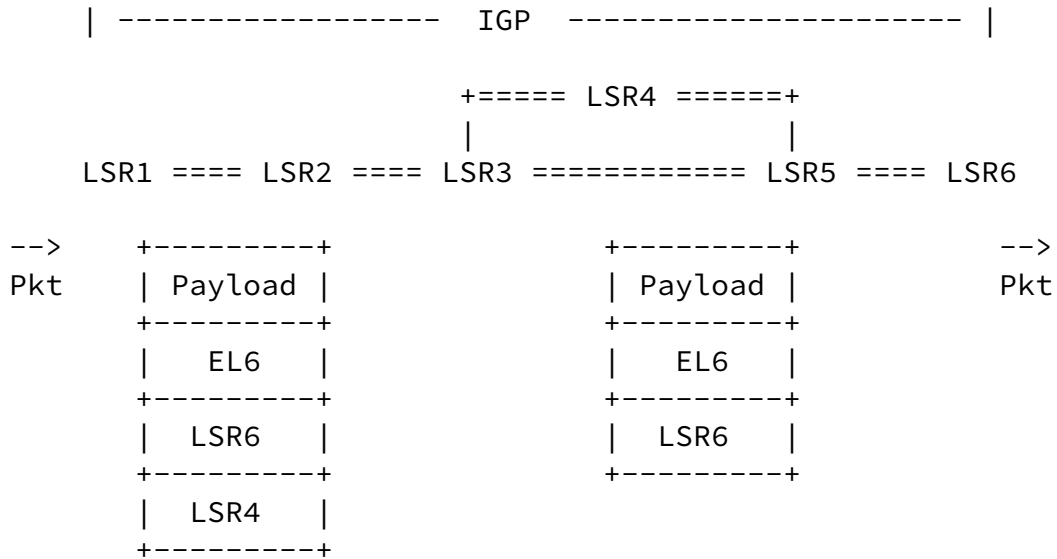


Figure 1: Single EL for TE tunnel LSR1->LSR4->LSR6

[4.2.](#) Example 2 Entropy Block advertisement through global configuration

In this example, Entropy Block is globally configured by NMS. All LSRs have the identical Entropy Block. One traffic flow is forwarded through a Traffic Engineering (TE) tunnel, LSR1->LSR4->LSR6. LSR2 and LSR3 have Less entropy reachability capability.

Considering the entropy reachability capability on some LSRs, more entropy labels are encapsulated in the MPLS stack. On ingress LSR, entropy value is calculated from packet received, and adjusted to Entropy Block globally configured.

On transit LSRs, such as LSR2 or LSR3, the entropy information in the packet is utilized for load balancing. When packet reaches LSR4, the steering point of the TE tunnel, the topmost tunnel label is popped and the first entropy label is popped. Then packet continues to egress LSR, LSR6. After the second tunnel label is popped, the second entropy label is popped and packet is sent out.

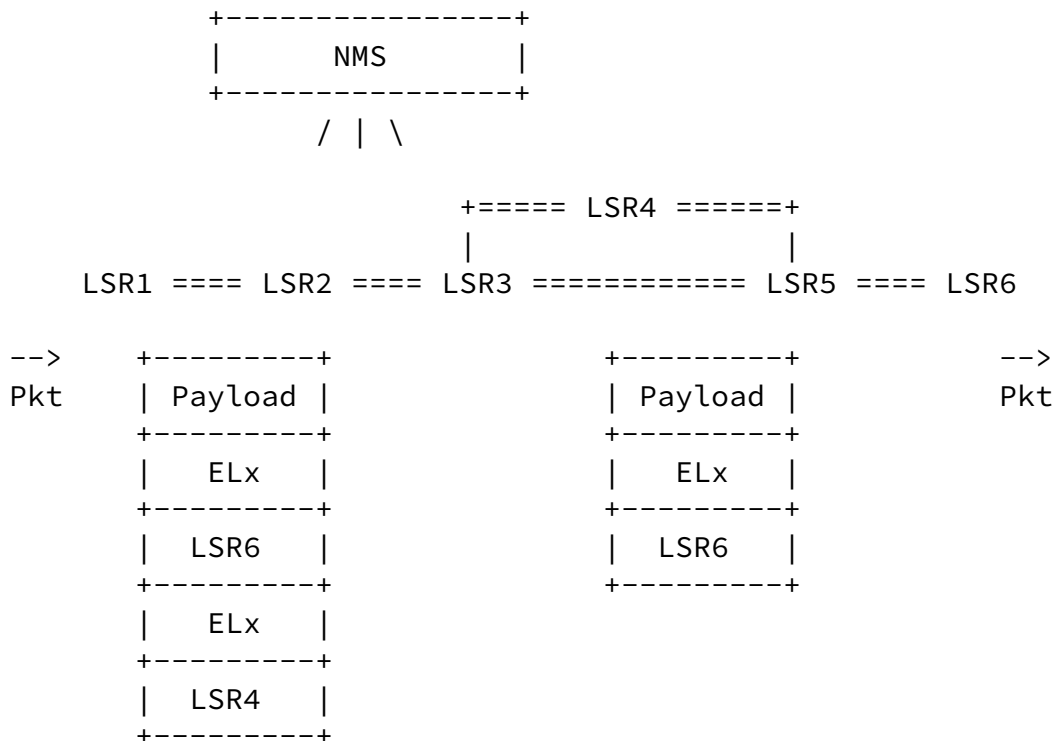


Figure 2: Multiple ELs for TE tunnel LSR1->LSR4->LSR6

[5.](#) Signaling for Entropy Block

The Entropy Block is advertised through signals (E.g. IGP) with new extensions. It could be represented as range size together with base label value, or the two boundary label values.

The detail formats of the extensions are left for further discussion.

[6.](#) Issues

[6.1.](#) Anycast Segment

In Segment Routing networks [[RFC8402](#)] when anycast segment is deployed, special care is required.

An anycast segment is not reference a particular node, but a group of nodes. In case each node in the group has different Entropy Block, ingress LSR cannot encapsulate an appropriate entropy label behind the anycast segment in the MPLS stack.

To ease the use of Entropy Block with anycast segment, it is recommended to configure identical Entropy Block on all nodes of a particular anycast group. Using an anycast segment without

configuring identical Entropy Block on all nodes belonging to the same anycast group may lead to misrouting.

[7.](#) IANA Considerations

TBD.

[8.](#) Security Considerations

With the use of the extensions defined in this document, Entropy information will be used to program the MPLS data plane [[RFC3031](#)]. Using an anycast segment without configuring identical Entropy Block on all nodes belonging to the same anycast group may lead to misrouting.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", [RFC 3031](#), DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.
- [RFC6391] Bryant, S., Ed., Filsfils, C., Drafz, U., Kompella, V., Regan, J., and S. Amante, "Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network", [RFC 6391](#), DOI 10.17487/RFC6391, November 2011, <<https://www.rfc-editor.org/info/rfc6391>>.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", [RFC 6790](#), DOI 10.17487/RFC6790, November 2012, <<https://www.rfc-editor.org/info/rfc6790>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

9.2. Informative References

- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L.,

Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

[RFC8662] Kini, S., Kompella, K., Sivabalan, S., Litkowski, S., Shakir, R., and J. Tantsura, "Entropy Label for Source Packet Routing in Networking (SPRING) Tunnels", [RFC 8662](#), DOI 10.17487/RFC8662, December 2019, <<https://www.rfc-editor.org/info/rfc8662>>.

Authors' Addresses

Jiang He (editor)
Ericsson
No. 5, Lize east street, Chaoyang district
Beijing 100102
CN

Email: jiang.he@ericsson.com

Bolin Nie
Ericsson
No. 5, Lize east street, Chaoyang district
Beijing 100102
CN

Email: zephyr.nie@ericsson.com

Zhenning Zhao
Ericsson
No. 5, Lize east street, Chaoyang district
Beijing 100102
CN

Email: navid.zhao@ericsson.com