

INTERNET-DRAFT
Intended Status: Standards Track
Expires: May 7, 2016

Luyuan Fang
Microsoft
Fabio Chiussi

Barak Gafni
Mellanox
Nabil Bitar
Alcatel Lucent
Zhenbin Li
Huawei
Rob Raszuk
Bloomberg
Andrew Qu
MediaTek

November 4, 2015

MPLS Label Forwarding with No Swapping
draft-fang-mpls-label-forwarding-no-swap-02

Abstract

This document defines MPLS label forwarding operation with no label swapping as a new MPLS label operation extension to the existing basic forwarding operation of label push, pop, and swap.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2015 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 (<http://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.

Table of Contents

1. Introduction	3
2. Terminology	4
3. Label Forwarding Operation as Defined in RFC3031	4
4. Label Forwarding with No Swap	5
5. Proposed Text to Update RFC3031	6
6. Security Considerations	7
7. IANA Considerations	7
8. References	7
8.1 Normative References	7
8.2 Informative References	7
Authors' Addresses	8

1. Introduction

MPLS forwarding operation as defined in [RFC3031] has three basic operations that must be performed on the labels at the network nodes: push, swap, and pop. This document describes an additional operation action: label forwarding with no swap. Currently, using the same label as both incoming and outgoing label is typically achieved by "swapping" the incoming label with an identical outgoing label. In order to improve processing efficiency and memory usage reduction, a simple label forwarding operation with no swap is desirable. This operation should be specified in [RFC3031]

When MPLS Architecture [RFC3031] was defined, the three types of label operation were sufficient, since labels were always only locally meaningful. Label swap operation is performed at a Label Switched Router (LSR) which is not an MPLS edge node, while label push and pop can be performed at an MPLS edge node for label imposition and deposition. Penultimate hop popping can also be performed at the penultimate hop for improved efficiency when appropriate. Since the labels are assigned independently in distributed fashion in a non-traffic engineered basic MPLS networks, it is not possible nor necessary to coordinate the label assignment. Therefore, the label swapping function is sufficient and effective for LSR.

With the increased interests and large scale development of Software-Defined Networking (SDN), several new use cases have emerged where domain-wide labels are used. For example, central controller assigned MPLS labels become one of the options for MPLS based forwarding. Using a single label is used to traverse multiple hops along the Label Switched Path (LSP) become desirable. Relevant examples of the use of domain-wide labels are described in [I-D.fang-mpls-hsdn-for-hsdc], [I-D.ietf-spring-segment-routing], and [I-D.li-mpls-global-label-usecases]

Domain-wide labels do not need to be swapped at the switches. Therefore, in order to allow efficient handling of domain-wide labels, there is a need to extend the label operation with one more action type - forwarding with no swap.

The performance and memory efficiency can be increased by performing simple forwarding function than swapping the labels with the identical identifier. This does not otherwise change the fundamentals of MPLS architecture and label encoding as defined by [RFC3031] [RFC3032].

2. Terminology

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].

- o Incoming Label Map (ILM): It maps each incoming label to a set of NHLFEs. It is used when forwarding packets that arrive as labeled packets.
- o Label forwarding: A simple forwarding paradigm allowing streamlined forwarding of data by using labels to identify classes of data packets which are treated indistinguishably when forwarding without label swapping.
- o Label swap: The basic existing forwarding operation consisting of looking up an incoming label to determine the outgoing label encapsulation, port, and other data handling information.
- o Label swapping: A forwarding paradigm allowing streamlined forwarding of data by using labels to identify classes of data packets which are treated indistinguishably when forwarding.
- o Label Switched Path (LSP): The path through one or more LSRs at one level of the hierarchy followed by a packets in a particular forwarding equivalence class (FEC).
- o Label Switching Router (LSR): An MPLS node which is capable of forwarding native L3 packets.
- o MPLS edge node: An MPLS node that connects an MPLS domain with a node which is outside of the domain, either because it does not run MPLS, and/or because it is in a different domain. Note that if an LSR has a neighboring host which is not running MPLS, that that LSR is an MPLS edge node.
- o NHLFE: Next Hop Label Forwarding Entry
- o Software-Defined Networking (SDN): an architecture that decouples the network control and forwarding functions to enable the network control to be directly programmable and the underlying infrastructure to be abstracted for applications and network services.

3. Label Forwarding Operation as Defined in RFC3031

Section 3.10 in [RFC3031] states the following:

"The "Next Hop Label Forwarding Entry" (NHLFE) is used when forwarding a labeled packet. It contains the following information:

1. the packet's next hop
2. the operation to perform on the packet's label stack; this is one of the following operations:
 - a) replace the label at the top of the label stack with a specified new label
 - b) pop the label stack
 - c) replace the label at the top of the label stack with a specified new label, and then push one or more specified new labels onto the label stack.

..."

Section 3.13 in [RFC3031] states the following:

"Label swapping is the use of the following procedures to forward a packet. In order to forward a labeled packet, a LSR examines the label at the top of the label stack. It uses the ILM to map this label to an NHLFE. Using the information in the NHLFE, it determines where to forward the packet, and performs an operation on the packet's label stack. It then encodes the new label stack into the packet, and forwards the result."

It is clear that, in order to forward a labeled packet, [RFC3031] mandates to perform one of three possible operations: swap, pop, or swap/push.

Therefore, in order to forward a labeled packet with a label that does not need to be swapped according to the standard, the label at the top of the label stack must be swapped with an identical label.

Simply forwarding the labeled packet without swapping the label does not conform with [RFC3031].

4. Label Forwarding with No Swap

Label forwarding is the use of the following procedures to forward a packet.

Same as in [RFC3031], in order to forward a labeled packet, a LSR examines the label at the top of the label stack. It uses the ILM to

map this label to an NHLFE. Using the information in the NHLFE, it determines where to forward the packet, and performs an operation on the packet's label stack.

Unlike in label swapping, label forwarding does not remove the incoming label and encodes the new label stack into the packet as in label swapping, it forwards the packet with the same label stack as the incoming stack, to the outgoing interface. Other processing may be involved in selecting the outgoing interface, for example, load balancing through IP deader hashing or use of Entropy label [RFC6790].

5. Proposed Text to Update RFC3031

In order to allow efficient processing of domain-wide labels, which do not need to be swapped, we propose to add a fourth possible operation "Forward No Swap" to [RFC3031].

Accordingly, [RFC3031] should be modified as follows.

1. the packet's next hop
 2. the operation to perform on the packet's label stack; this is one of the following operations:
 - a) replace the label at the top of the label stack with a specified new label
 - b) pop the label stack
 - c) replace the label at the top of the label stack with a specified new label, and then push one or more specified new labels onto the label stack
 - d) keep the label at the top of the label stack unchanged.
- ...

6. Security Considerations

The MPLS label forwarding operation specified herein does not raise any security issues that are not already present in either the MPLS architecture [RFC3031] or in MPLS label encoding [RFC3032].

In addition, general MPLS and GMPLS considerations and MPLS security defense techniques are documented in [RFC5920].

7. IANA Considerations

None.

8. References

8.1 Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [RFC6790] Kompella, K., et.al, "The Use of Entropy Labels in MPLS Forwarding", November 2012.

8.2 Informative References

- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
- [I-D.fang-mpls-hsdn-for-hsdc] L. Fang, et al., "MPLS-Based Hierarchical SDN for Hyper-Scale DC/Cloud", draft-fang-mpls-hsdn-for-hsdc-04 (work in progress), July 2015.
- [I-D.ietf-spring-segment-routing] Filsfils, C. et al., "Segment Routing Architecture", draft-ietf-spring-segment-routing-06 (work in progress), October 2015.
- [I-D.li-mpls-global-label-usecases] Z. Li, et al., "Usecases of MPLS Global Label, draft-li-mpls-global-label-usecases-03 (work in progress), October 2015.

Authors' Addresses

Luyuan Fang
Microsoft
5600 148th Ave NE
Redmond, WA 98052
Email: lufang@microsoft.com

Fabio Chiusi
Seattle, WA 98116
Email: fabiochiussi@gmail.com

Barak Gafni
Mellanox
6 Habarzel St.
Tel Aviv, Israel
Email: gbarak@mellanox.com

Nabil Bitar
Alcatel Lucent

Zhenbin Li
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Robert Raszuk
Email: robert@raszuk.net

Andrew Qu
MediaTek
2860 Junction Ave.
San Jose, CA 95134
Email: andrew.qu@mediatek.com