

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: August 16, 2014

M. Chen  
X. Xu  
Z. Li  
Huawei  
L. Fang  
Microsoft  
G. Mirsky  
Ericsson  
February 12, 2014

MultiProtocol Label Switching (MPLS) Source Label  
draft-chen-mpls-source-label-02

Abstract

An MultiProtocol Label Switching (MPLS) label is originally defined to identify a Forwarding Equivalence Class (FEC), a packet is assigned to a specific FEC based on its network layer destination address. It's difficult or even impossible to derive the source information from the label. For some applications, source identification is a critical requirement. For example, performance monitoring, traffic matrix measurement and collection, where the monitoring node needs to identify where a packet was sent from.

This document introduces the concept of Source Label (SL) that is carried in the label stack and used to identify the ingress Label Switching Router (LSR) of an Label Switched Path (LSP).

Requirements Language

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].

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 <http://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 August 16, 2014.

#### Copyright Notice

Copyright (c) 2014 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. Problem Statement and Introduction . . . . .	3
2. Source Label . . . . .	4
3. Use Cases . . . . .	4
3.1. Performance Measurement . . . . .	4
3.2. Traffic Matrix Measurement and Steering . . . . .	5
3.3. Source Filtering . . . . .	7
4. Data Plane Processing . . . . .	7
4.1. Ingress LSR . . . . .	7
4.2. Transit LSR . . . . .	8
4.3. Egress LSR . . . . .	8
4.4. Penultimate Hop LSR . . . . .	8
5. Source Label Signaling . . . . .	8
5.1. Source Label Capability Signaling . . . . .	8
5.1.1. LDP Extensions . . . . .	8
5.1.2. BGP Extensions . . . . .	9
5.1.3. RSVP-TE Extensions . . . . .	10
5.2. Source Label Distribution . . . . .	10
6. IANA Considerations . . . . .	11
6.1. Source Label Indication . . . . .	11
6.2. LDP Source Label Capability TLV . . . . .	11
6.3. BGP Source Label Capability Attribute . . . . .	11
6.4. RSVP-TE Source Label Capability . . . . .	11
7. Security Considerations . . . . .	11
8. Acknowledgements . . . . .	12
9. References . . . . .	12
9.1. Normative References . . . . .	12

9.2. Informative References . . . . .	12
Authors' Addresses . . . . .	13

## 1. Problem Statement and Introduction

An MultiProtocol Label Switching (MPLS) label [RFC3031] is originally defined for packet forwarding and assumes the forwarding/destination address semantics. As no source address information is carried in the label stack, there is no way to directly derive the source address information from the label or label stack.

MPLS LSPs can be categorized into four different types:

Point-to-Point (P2P)

Point-to-Multipoint (P2MP)

Multipoint-to-Point (MP2P)

Multipoint-to-Multipoint (MP2MP)

For Resource Reservation Protocol Traffic Engineering (RSVP-TE) [RFC3209] based P2P and P2MP LSPs, the source address information may be implicitly derived from the label when Penultimate Hop Popping (PHP) is disabled. Note that such LSP may be characterized as MPLS-TP LSP [RFC5960]. But it requires that some further information is used (e.g., control plane information).

For Label Distribution Protocol (LDP) based LSPs [RFC5036] [RFC6388], Layer 3 Private Network (L3VPN) and Virtual Local Area Network (VPLS) LSPs that normally belong to P2MP, MP2P and MP2MP LSPs, ingress LSR that sent particular MPLS frame over P2MP, MP2P or MP2MP LSP cannot be identified by egress LSR.

Comparing to the pure IP forwarding where both source and destination addresses are encoded in the IP packet header, the essential issue of the MPLS encoding is that the label stack does not explicitly include any source address information, i.e., a Source Label (SL). For some applications, source identification is a critical requirement. For example, performance monitoring, the monitoring nodes need to identify where packets were sent from and then can count the packets according to some constraints. In addition, traffic matrix measurement and collection is the precondition of traffic steering, and capable of traffic steering is an important requirement of Software Defined Network (SDN). To measure and collect traffic matrix information, the source address information is necessary.

In addition, Segment Routing [I-D.filsfils-rtgwg-segment-routing] also explicitly points out that there are requirements to preserve the ingress information to fulfill the accounting and billing purposes.

This document introduces the concept of Source Label. An SL uniquely identifies a node within an administrative domain, it is carried in the label stack and used to identify one of the ingress LSR(s) of an LSP.

## 2. Source Label

A Source Label is defined to uniquely identify a node that is (one of) the ingress LSR(s) to a specific LSP. In its function as a Source Label, it MUST be unique within a domain. In cases where a Source Label is used across domains it MUST be unique within the scope it is used.

Source Labels SHOULD NOT be used for forwarding. The Source Labels are allocated from a dedicated label space that is completely different from the space of the normal Forwarding Labels. Configuration system (e.g., static configuration) is one way to make sure the uniqueness of each SL assigned to specific LSR. There may be some other potential dynamic solutions that can be used for SL allocation and distribution. This is out of the scope of this document.

In order to indicate whether a label is a source label, a Source Label Indicator (SLI) is introduced. The SLI is a (extended) special purpose label that is placed immediately before the source label in the label stack, which is used to indicate that the next label in the label stack is a source label. The value of SLI is TBD1.

## 3. Use Cases

This section outlines a number of use cases where solutions built on Source Label.

### 3.1. Performance Measurement

There are two typical types of performance measurement: one is active performance measurement, and the other is passive performance measurement.

In active performance measurement the receiver measures the injected packets to evaluate the performance of a path. The active measurement measures the performance of the extra injected packets. The IP Performance Metrics (IPPM) working group has defined

specifications [RFC4656][RFC5357] for the active performance measurement.

In passive performance measurement, no artificial traffic is injected into the flow and measurements are taken to record the performance metrics of the real traffic. The Multiprotocol Label Switching (MPLS) PM protocol [RFC6374] for packet loss is an example of passive performance measurement, but it can only apply to MPLS-TE LSPs. For a specific receiver, in order to count the received packets of a flow, it has to know whether a received packet belongs to which target flow under test and the source identification is a critical condition.

As discussed in the previous section, the existing MPLS label or label stack do not carry the source information. So, for an LSP, the ingress LSR can put a source label in the label stack, and then the egress LSR can use the source label for packets identifying and counting.

### 3.2. Traffic Matrix Measurement and Steering

A Traffic Matrix (TM) provides, for every ingress node (i) into the network and every egress node (j) out of the network, the volume of traffic  $T(i,j)$  from i to j over a given time interval.

Since the ingress node knows the source and destination of the traffic, it's normal to measure the traffic matrix at every ingress node. But in some scenarios, it may need to measure the traffic at the egress or intermediate nodes. Taking Figure 1 as an example, from the west to east point of view, there are three ingress nodes (I1, I2 and I3) and three egress nodes (E1, E2 and E3), A, B and C are intermediate nodes. It is not necessary to measure the traffic matrix of the whole network all the time, it sometimes just wants to know the received traffic matrix of a specific egress node (e.g., E2). So, to measure received traffic matrix at node E2 would be then a better choice.




Figure 1: Traffic Matrix Measurement and Steering

In addition, for an intermediate node (e.g., node A), to steer traffic from congested path (e.g., path A-C) to idle path (e.g., path A-B), it needs to identify which flows contribute to the congestion and then determine which flows (e.g., the flows from specific ingress node) should be moved to the idle path.

Another scenario is domain exit traffic steering. Taking figure 2 as an example, node D is the domain gateway and has multiple exit links. Sometime, it may need to perform ingress/source node based traffic steering. It means that traffic from specific ingress node is required to be forwarded through specific exit link. For example, traffic from node A is required to be sent along with link 1, traffic from node B is required to be sent along with link 2, and traffic from node C is required to be sent along with link 3. To achieve this, node D needs to identify from which a flow is sent.




Figure 2: Domain Exit Traffic Steering

According above, wherever at egress or intermediate node, source identification is necessary. It should be possible to configure the ingress LSR to put the source label into the label stack to enable the egress and intermediate LSR to identify, measure and steer the traffic.

### 3.3. Source Filtering

Network Ingress Filtering [RFC2827] is an important tool to defeat DoS attacks and is widely deployed. In the past, since there is no source information carried in the stack, it's impossible to perform source filtering. With the Source Label, it enables to filter the packets with specific Source Label.

## 4. Data Plane Processing

### 4.1. Ingress LSR

For an LSP, the ingress LSR MUST make sure that the egress LSR is able to process the Source Label before inserting an SL and SLI into the label stack. Therefore, an egress LSR SHOULD signal (see Section 5.1) to the ingress LSR whether it is able to process the Source Label. Once the ingress LSR knows that the egress LSR can process Source Label, it can choose whether or not to insert the SL and SLI into the label stack.

When an SL to be included in a label stack, the steps are as follows:

1. Push the SL label, the BoS bit for the SL depends on whether the SL is the bottom label;
2. Push the SLI, the TTL and TC field for the SLI SHOULD be set to the same values as for the LSP Label (L);

### 3. Push the LSP Label (L) .

Then the label stack looks like: <...L, SLI, SL...>. There may be multiple pairs of SLI and SL inserted into the label stack, each pair is related to an LSP. For the given LSP, only one pair of SLI and SL SHOULD be inserted.

### 4.2. Transit LSR

There is no change in forwarding behavior for transit LSRs. But if a transit LSR can recognize the SLI, it can use the SL to collect traffic throughput and/or measure the performance of the LSP.

### 4.3. Egress LSR

When an egress LSR receives a packet with a SLI/SL pair, if the egress LSR is able to process the SL; it pops the LSP label (if any), SLI and SL; then processes remaining packet header as normal. If the egress LSR is not able to process the SL, the packet SHOULD be dropped as specified for the handling of any unknown label according to [RFC3031].

### 4.4. Penultimate Hop LSR

There is no change in forwarding behavior for the penultimate hop LSR.

## 5. Source Label Signaling

Source label signaling includes two aspects: one is source label capability signaling, the other is source label distribution.

### 5.1. Source Label Capability Signaling

Before inserting a source label in the label stack, an ingress LSR MUST know whether the egress LSR is able to process the source label. Therefore, an egress LSR should signal to the ingress LSRs its ability to process the Source Label. This is called Source Label Capability (SLC), it is very similar to the "Entropy Label Capability (ELC)" [RFC6790].

#### 5.1.1. LDP Extensions

A new LDP TLV [RFC5036], SLC TLV, is defined to signal an egress's ability to process source label. The SLC TLV may appear as an Optional Parameter of the Label Mapping Message. The presence of the SLC TLV in a Label Mapping Message indicates to ingress LSRs that the egress LSR can process source labels for the associated LSP.



The structure of the SLC TLV is shown below.




Figure 1: Source Label Capability TLV

This U bit MUST be set to 1. If the SLC TLV is not understood by the receiver, then it MUST be ignored.

This F bit MUST be set to 1. Since the SLC TLV is going to be propagated hop-by-hop, it should be forwarded even by nodes that may not understand it.

Type: TBD2.

Length field: This field specifies the total length in octets of the SLC TLV and is defined to be 0.

An LSR that receives a Label Mapping with the SLC TLV but does not understand it MUST propagate it intact to its neighbors and MUST NOT send a notification to the sender (following the meaning of the U- and F-bits). If the LSR has no other neighbors and does not understand the SLC TLV, means it is the ingress LSR, it could just ignore it. An LSR X may receive multiple Label Mappings for a given FEC F from its neighbors. In its turn, X may advertise a Label Mapping for F to its neighbors. If X understands the SLC TLV, and if any of the advertisements it received for FEC F does not include the SLC TLV, X MUST NOT include the SLC TLV in its own advertisements of F. If all the advertised Mappings for F include the SLC TLV, then X MUST advertise its Mapping for F with the SLC TLV. If any of X's neighbors resends its Mapping, sends a new Mapping or sends a Label Withdraw for a previously advertised Mapping for F, X MUST re-evaluate the status of SLC for FEC F, and, if there is a change, X MUST re-advertise its Mapping for F with the updated status of SLC.

#### 5.1.2. BGP Extensions

When Border Gateway Protocol (BGP) [RFC4271] is used for distributing Network Layer Reachability Information (NLRI) as described in, for example, [RFC3107], [RFC4364], the BGP UPDATE message may include the SLC attribute as part of the Path Attributes. This is an optional, transitive BGP attribute of value TBD3. The inclusion of this attribute with an NLRI indicates that the advertising BGP router can process source labels as an egress LSR for all routes in that NLRI.

A BGP speaker S that originates an UPDATE should include the SLC attribute only if both of the following are true:

A1: S sets the BGP NEXT\_HOP attribute to itself AND

A2: S can process source labels.

Suppose a BGP speaker T receives an UPDATE U with the SLC attribute. T has two choices. T can simply re-advertise U with the SLC attribute if either of the following is true:

B1: T does not change the NEXT\_HOP attribute OR

B2: T simply swaps labels without popping the entire label stack and processing the payload below.

An example of the use of B1 is Route Reflectors. However, if T changes the NEXT\_HOP attribute for U and in the data plane pops the entire label stack to process the payload, T MAY include an SLC attribute for UPDATE U' if both of the following are true:

C1: T sets the NEXT\_HOP attribute of U' to itself AND

C2: T can process source labels. Otherwise, T MUST remove the SLC attribute.

#### 5.1.3. RSVP-TE Extensions

[RFC5420] introduces the LSP\_ATTRIBUTES object, it gives a perfect way to carry LSP attribute through the object. To signal the Source Label Capability in RSVP-TE [RFC3209], this document defines a flag in the Attribute Flags TLV of the the LSP\_ATTRIBUTES object [RFC3209].

The presence of the SLC flag in a Path message indicates that the ingress can process source labels in the upstream direction; this only makes sense for a bidirectional LSP and MUST be ignored otherwise. The presence of the SLC flag in a Resv message indicates that the egress can process source labels in the downstream direction. The bit number for the SLC flag is TBD4.

#### 5.2. Source Label Distribution

Based on the Source Label, an egress or intermediate LSR can identify from where an MPLS packet is sent. To achieve this, the egress and/or intermediate LSRs have to know which ingress LSR is related to which Source Label before using the Source Label to derive the source information. Therefore, there needs to be a mechanism to distribute

the mapping information between an ingress LSR and its Source Label. This can be done, for example, by defining extensions to LDP, BGP, RSVP-TE and/or Interior Gateway Protocol (IGP) to distribute to source label mapping. The source label distribution will be defined in another document(s).

## 6. IANA Considerations

### 6.1. Source Label Indication

IANA is required to allocate a special purpose label (TBD1) for the Source Label Indicator (SLI) from the "Multiprotocol Label Switching Architecture (MPLS) Label Values" Registry.

### 6.2. LDP Source Label Capability TLV

IANA is requested to allocate a value of TBD2 from the IETF Consensus range (0x0001-0x07FF) in the "TLV Type Name Space" registry as the "Source Label Capability TLV".

### 6.3. BGP Source Label Capability Attribute

IANA is requested to allocate a Path Attribute Type Code TBD3 from the "BGP Path Attributes" registry as the "BGP Source Label Capability Attribute".

### 6.4. RSVP-TE Source Label Capability

IANA is requested to allocate a new bit from the "Attribute Flags" sub-registry of the "Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Parameters" registry.

Bit	Name	Attribute	Attribute	RRO
No		Flags Path	Flags Resv	
-----+-----+-----+-----+-----				
TBD4	Source Label Capability	Yes	Yes	No

## 7. Security Considerations

This document does not introduce extra security issues. On the contrary, with the Source Label carried in the stack, it may bring additional security enhancement that enables an LSR to perform source label based checking and/or filtering.

## 8. Acknowledgements

The process of "Source Label Capability Signaling" is largely referred to the process of "ELC signaling"[RFC6790].

The authors would like to thank Carlos Pignataro, Loa Andersson for their review, suggestion and comments to this document.

## 9. References

### 9.1. 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.
- [RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", RFC 3107, May 2001.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", RFC 5036, October 2007.
- [RFC5420] Farrel, A., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.
- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011.

### 9.2. Informative References

- [I-D.filsfils-rtgwg-segment-routing]  
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., and E. Crabbe, "Segment Routing Architecture", draft-filsfils-rtgwg-segment-routing-01 (work in progress), October 2013.

- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006.
- [RFC4761] Kompella, K. and Y. Rekhter, "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.
- [RFC5960] Frost, D., Bryant, S., and M. Bocci, "MPLS Transport Profile Data Plane Architecture", RFC 5960, August 2010.
- [RFC6388] Wijnands, IJ., Minei, I., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6388, November 2011.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, November 2012.

#### Authors' Addresses

Mach(Guoyi) Chen  
Huawei

Email: mach.chen@huawei.com

Xiaohu Xu  
Huawei

Email: xuxiaohu@huawei.com

Zhenbin Li  
Huawei

Email: lizhenbin@huawei.com

Luyuan Fang  
Microsoft

Email: lufang@microsoft.com

Greg Mirsky  
Ericsson

Email: Gregory.mirsky@ericsson.com