

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 15, 2017

X. Xu, Ed.
Huawei
R. Raszuk
Bloomberg LP
U. Chunduri

L. Contreras
Telefonica I+D
L. Jalil
Verizon
October 12, 2016

Connecting MPLS-SPRING Islands over IP Networks
draft-xu-mpls-spring-islands-connection-over-ip-00

Abstract

MPLS-SPRING is an MPLS-based source routing paradigm in which a sender of a packet is allowed to partially or completely specify the route the packet takes through the network by imposing stacked MPLS labels to the packet. To facilitate the incremental deployment of this new technology, this document describes a mechanism which allows the outermost LSP be replaced by an IP-based tunnel.

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 April 15, 2017.

Copyright Notice

Copyright (c) 2016 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	2
1.1. Requirements Language	3
2. Terminology	3
3. Packet Forwarding Procedures	3
4. Acknowledgements	4
5. IANA Considerations	4
6. Security Considerations	4
7. References	4
7.1. Normative References	4
7.2. Informative References	4
Authors' Addresses	6

1. Introduction

MPLS-SPRING [I-D.ietf-spring-segment-routing-mpls] is a MPLS-based source routing paradigm in which a sender of a packet is allowed to partially or completely specify the route the packet takes through the network by imposing stacked MPLS labels to the packet. To facilitate the incremental deployment of this new technology, this document describes a mechanism which allows the outermost LSP to be replaced by an IP-based tunnel (e.g., MPLS-in-IP/GRE tunnel [RFC4023], MPLS-in-UDP tunnel [RFC7510] or MPLS-in-L2TPv3 tunnel [RFC4817] and etc) when the nexthop along the LSP is not MPLS-SPRING-enabled. The tunnel destination address would be the address of the egress of the outmost LSP (e.g., the egress of the active node segment).

This mechanism is much useful in the MPLS-SPRING-based Service Function Chaining (SFC) case [I-D.xu-sfc-using-mpls-spring] where only a few specific routers (e.g., Service Function Forwarders (SFF) and classifiers) are required to be MPLS-SPRING-capable while the remaining routers are just required to support IP forwarding capability. In addition, this mechanism is also useful in some specific Traffic Engineering scenarios where only a few routers (e.g., the entry and exit nodes of each plane in the dual-plane network) are specified as segments of explicit paths. In this way,

only a few routers are required to support the MPLS-SPRING capability while all the other routers just need to support IP forwarding capability, which would significantly reduce the deployment cost of this new technology. Furthermore, since there is no need to run any other label distribution protocol (e.g., LDP), the network provisioning is greatly simplified, which is one of the major claimed benefits of the MPLS-SPRING technology.

1.1. 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].

2. Terminology

This memo makes use of the terms defined in [RFC3031] and [I-D.ietf-spring-segment-routing-mpls].

3. Packet Forwarding Procedures

Assume an MPLS-SPRING-enabled router X prepares to forward an MPLS packet to the next node segment (i.e., the node segment of MPLS-SPRING-enabled router Y) which is identified by the top label of the MPLS packet. If the next-hop router of the best path to Y is a non-MPLS router, X couldn't map the packet's top label into a Next Hop Label Forwarding Entry (NHLFE), even though the top label itself is a valid incoming label. According to the following specification as quoted from Section 3.22 of [RFC3031], the MPLS packet would be discarded in the current MPLS implementations:

"When a labeled packet is traveling along an LSP, it may occasionally happen that it reaches an LSR at which the ILM does not map the packet's incoming label into an NHLFE, even though the incoming label is itself valid...Unless it can be determined (through some means outside the scope of this document) that neither of these situations obtains, the only safe procedure is to discard the packet. "

This document proposes an improved procedure to deal with the above case. The basic idea is to set an IP tunnel towards the egress of topmost LSP as the NHLFE of that incoming label. More specifically, if the label is not a Penultimate Hop Popping (PHP) label (i.e., the NP-flag [I-D.ietf-isis-segment-routing-extensions] associated with the corresponding prefix SID of that top label is set), X SHOULD swap the label to the corresponding label significant to Y and then encapsulate the MPLS packet into the IP-based tunnel towards Y. The tunnel destination address is the IP address of Y (e.g., the /32 or

/128 prefix FEC associated with that top label) and the tunnel source address is the IP address of X. If the label is a PHP label and not at the bottom of the label stack, X SHOULD pop that label before performing the above MPLS over IP encapsulation. The IP encapsulated MPLS packet would be forwarded according to the IP routing table. Upon receipt of that IP encapsulated MPLS packet, Y would decapsulate it and then process the decapsulated MPLS packet accordingly. As for which tunnel encapsulation type should be used by X, it can be manually specified on X or be learnt from Y's advertisement of its tunnel encapsulation capability. How to advertise the tunnel encapsulation capability using IS-IS or OSPF are specified in [I-D.xu-isis-encapsulation-cap] and [I-D.ietf-ospf-encapsulation-cap] respectively.

4. Acknowledgements

Thanks Joel Halpern, Bruno Decraene and Loa Andersson for their insightful comments on this draft.

5. IANA Considerations

No IANA action is required.

6. Security Considerations

TBD.

7. References

7.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.

7.2. Informative References

[I-D.ietf-isis-segment-routing-extensions]
Previdi, S., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., Decraene, B., and J. Tantsura, "IS-IS Extensions for Segment Routing", draft-ietf-isis-segment-routing-extensions-07 (work in progress), June 2016.

- [I-D.ietf-ospf-encapsulation-cap]
Xu, X., Decraene, B., Raszuk, R., Chunduri, U., Contreras, L., and L. Jalil, "Advertising Tunnelling Capability in OSPF", draft-ietf-ospf-encapsulation-cap-00 (work in progress), October 2015.
- [I-D.ietf-spring-segment-routing-mpls]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Shakir, R., jefftant@gmail.com, j., and E. Crabbe, "Segment Routing with MPLS data plane", draft-ietf-spring-segment-routing-mpls-05 (work in progress), July 2016.
- [I-D.xu-isis-encapsulation-cap]
Xu, X., Decraene, B., Raszuk, R., Chunduri, U., Contreras, L., and L. Jalil, "Advertising Tunnelling Capability in IS-IS", draft-xu-isis-encapsulation-cap-06 (work in progress), November 2015.
- [I-D.xu-sfc-using-mpls-spring]
Xu, X., Shah, H., Contreras, L., and d. daniel.bernier@bell.ca, "Service Function Chaining Using MPLS-SPRING", draft-xu-sfc-using-mpls-spring-06 (work in progress), July 2016.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<http://www.rfc-editor.org/info/rfc3031>>.
- [RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, DOI 10.17487/RFC4023, March 2005, <<http://www.rfc-editor.org/info/rfc4023>>.
- [RFC4817] Townsley, M., Pignataro, C., Wainner, S., Seely, T., and J. Young, "Encapsulation of MPLS over Layer 2 Tunneling Protocol Version 3", RFC 4817, DOI 10.17487/RFC4817, March 2007, <<http://www.rfc-editor.org/info/rfc4817>>.
- [RFC7510] Xu, X., Sheth, N., Yong, L., Callon, R., and D. Black, "Encapsulating MPLS in UDP", RFC 7510, DOI 10.17487/RFC7510, April 2015, <<http://www.rfc-editor.org/info/rfc7510>>.

Authors' Addresses

Xiaohu Xu (editor)
Huawei

Email: xuxiaohu@huawei.com

Robert Raszuk
Bloomberg LP

Email: robert@raszuk.net

Uma Chunduri

Email: uma.chunduri@gmail.com

Luis M. Contreras
Telefonica I+D

Email: luismiguel.contrerasmurillo@telefonica.com

Luay Jalil
Verizon

Email: luay.jalil@verizon.com