Network Working Group
Internet-Draft
Intended status: Informational
Expires: February 5, 2015

X. Xu
Huawei
S. Sivabalan
Cisco
R. Raszuk
Individual
U. Chunduri
Ericsson
V. Lopezalvarez
Telefonica

August 4, 2014

Connecting MPLS-SPRING Islands over IP Networks
draft-xu-spring-islands-connection-over-ip-01

Abstract

MPLS-SPRING is a 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 using stacked MPLS labels. The current MPLS-SRPING architecture requires an end-to-end MPLS Label Switched Path (LSP) between any two MPLS-SPRING-enabled routers (e.g., two adjacent hops of a given explicit path). In order to enable MPLS-SPRING-enabled routers to be deployed even when there are non-MPLS routers along the path between two MPLS-SPRING-enabled routers, it is desirable to have an alternative, which allows the use of IP-based tunnels (e.g., GRE tunnels) to connect two MPLS-SPRING-enabled routers. This document describes a mechanism for such usage.

Table of Contents

1. Introduction

MPLS-SPRING [I-D.gredler-spring-mpls] [I-D.filsfils-spring-segment-routing-mpls] is a 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 using stacked MPLS labels. The current MPLS-SRPING architecture requires an end-to-end MPLS Label Switched Path (LSP) between any two MPLS-SPRING-enabled routers (e.g., two adjacent hops of a given explicit path). In order to enable MPLS-SPRING-enabled routers to be deployed even when there are non-MPLS routers along the path between two MPLS-SPRING-enabled routers, it is desirable to have an alternative, which allows the use of IP-based tunnels (e.g., GRE tunnels [RFC4023] or UDP tunnels [I-D.ietf-mpls-in-udp]) to connect two MPLS-SPRING-enabled routers which are specified as adjacent hops of a given explicit path. The tunnel destination address would be the address of next-hop MPLS-SPRING-enabled router along the explicit path, and this would cause the packet to be delivered to the next explicit hop. In this procedure, the ingress and egress of the IP-

based tunnel MUST support MPLS-SPRING features including the MPLS forwarding capability, whereas those transit routers along the path between them don't need to support any MPLS-SPRING features including the MPLS forwarding capability. The above mechanism is much useful for incrementally deployment of the MPLS-SPRING technology, especially in the MPLS-SPRING-based Service Function Chainning (SFC) case where only a few specific routers (e.g., service nodes and classifiers) are actually required to be MPLS-SPRING-capable.

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 [I-D.gredler-spring-mpls] and [I-D.filsfils-spring-segment-routing-mpls].

3. Packet Forwarding Procedures

Assume an MPLS-SPRING-enabled router X prepares to forward an MPLS packet to the next explicit hop Y which is identified by the top label of the MPLS packet, if the next-hop router Z which is physically adjacent to X is a non-MPLS-SPRING router, X would encapsulate the MPLS packet into an IP-based tunnel (e.g., GRE tunnel or UDP tunnel) where the tunnel destination is an IP address of Y (i.e., the /32 or /128 IGP prefix FEC corresponding to that top label) and the tunnel source is an IP address of X. If the top label is a Penultimate Hop Popping (PHP) label, X SHOULD pop that top label before performing the encapsulation. The IP encapsulated packet would be forwarded according to the IP forwarding table. Upon receipt of that IP encapsulated 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 dynamically learnt from Y's advertisement of its tunnel encapsulation capability. How to advertise the tunnel encapsulation capability is outside of the scope of this document.

4. Acknowledgements

Thanks Joel Halpern for his insightful comments on this draft.

5. IANA Considerations

No action is required for IANA.

6. Security Considerations

TBD.

7. References

7.1. Normative References

[I-D.filsfils-spring-segment-routing-mpls]
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 with MPLS data plane", draft-filsfils-spring-segment-routing-mpls-03 (work in progress), August 2014.

[I-D.gredler-spring-mpls]
Gredler, H., Rekhter, Y., Jalil, L., Kini, S., and X. Xu, "Supporting Source/Explicitly Routed Tunnels via Stacked LSPs", draft-gredler-spring-mpls-06 (work in progress), May 2014.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2. Informative References

[I-D.ietf-mpls-in-udp]
Xu, X., Sheth, N., Yong, L., Pignataro, C., and F. Yongbing, "Encapsulating MPLS in UDP", draft-ietf-mpls-in-udp-05 (work in progress), January 2014.

[RFC4023] Worster, T., Rekhter, Y., and E. Rosen, "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, March 2005.

Authors' Addresses

Xiaohu Xu
Huawei

Email: xuxiaohu@huawei.com

Siva Sivabalan
Cisco

Email: msiva@cisco.com

Robert Raszuk
Individual

Email: robert@raszuk.net

Uma Chunduri
Ericsson

Email: uma.chunduri@ericsson.com

Victor Lopezalvarez
Telefonica

Email: victor.lopezalvarez@telefonica.com