

idr
Internet-Draft
Intended status: Standards Track
Expires: January 8, 2022

Z. Zhang
Juniper Networks
July 07, 2021

MPLS Label Stacks in Tunnel Encapsulation Attribute
draft-zzhang-idr-tunnel-encapsulation-label-stack-00

Abstract

[RFC9012] defines an MPLS Label Stack sub-TLV for Tunnel Encapsulation Attribute, and specifies that it is to be pushed BEFORE other labels. This document clarifies the use case for that, and defines a new Tunnel Label Stack sub-TLV for a label stack to be pushed AFTER other labels (e.g., the label embedded in the NLRI for a labeled address family, and/or the stack in an MPLS Label Stack sub-TLV).

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 8, 2022.

Copyright Notice

Copyright (c) 2021 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.

1. Introduction

[RFC9012] defines an MPLS Label Stack sub-TLV for Tunnel Encapsulation Attribute and specifies that:

"If a packet is to be sent through the tunnel identified in a particular TLV, and if that TLV contains an MPLS Label Stack sub-TLV, then the label stack appearing in the sub-TLV MUST be pushed onto the packet before any other labels are pushed onto the packet."

Specifically, the label stack in the sub-TLV is to be pushed BEFORE any other labels are pushed. This may sound counter-intuitive, in that if a label stack (e.g. for Segment Routing) is to be used to steer traffic to the tunnel endpoint, the stack should be pushed AFTER other labels (e.g. the label embedded in the NLRI).

This document clarifies that it is NOT for steering traffic to but for steering AFTER the tunnel endpoint. Consider the following use case:

```

      controller
      .         .
      .         .
site1 --- PE1 ----- PE2 --- site2
```

Two sites are connected to two PEs respectively, and traffic steering is desired within each site not just among the PEs. While PE2 could push the label stack used for steering within site2, there may be situations where PE2 is not a device capable of pushing a large label stack so PE1 is tasked with pushing the label stack that is used after the tunnel endpoint PE2, within site2.

2. Tunnel Label Stack sub-TLV

Notice that in the above example, the route update that PE1 receives could be from the controller instead of PE2. The controller may want to steer traffic both within site2 and between PE1 and PE2, yet [RFC9012](#) currently does not specify how to signal the label stack used to reach the tunnel endpoint.

To be able to signal that, this document defines a new Tunnel Label Stack sub-TLV. It has exactly the same syntax as the existing MPLS Label Stack sub-TLV, but with a semantics that it is pushed AFTER other labels are pushed.

[3.](#) Security Considerations

This document does not introduce any new security issues besides what is already discussed in [RFC9012](#).

[4.](#) IANA Assignments

IANA is requested to assign a new sub-TLV type for "Tunnel Label Stack" from "BGP Tunnel Encapsulation Attribute Sub-TLVs" registry, in the 0~127 range.

[5.](#) Acknowledgements

The authors thank Eric Rosen and John Scudder for their valuable comments and suggestions.

[6.](#) Normative References

- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", [RFC 9012](#), DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.

Author's Address

Zhaohui Zhang
Juniper Networks

Email: zzhang@juniper.net

Zhang

Expires January 8, 2022

[Page 3]