

Internet Engineering Task Force (IETF)
Request for Comments: 7213
Category: Standards Track
ISSN: 2070-1721

D. Frost
Blue Sun
S. Bryant
Cisco Systems
M. Bocci
Alcatel-Lucent
June 2014

MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing

Abstract

The MPLS Transport Profile (MPLS-TP) is the set of MPLS protocol functions applicable to the construction and operation of packet-switched transport networks. This document presents considerations for link-layer addressing of Ethernet frames carrying MPLS-TP packets.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7213>.

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.

1. Introduction

The MPLS Transport Profile (MPLS-TP) [RFC5921] is the set of protocol functions that meet the requirements [RFC5654] for the application of MPLS to the construction and operation of packet-switched transport networks. The MPLS-TP data plane consists of those MPLS-TP functions concerned with the encapsulation and forwarding of MPLS-TP packets and is described in [RFC5960].

This document presents considerations for link-layer addressing of Ethernet frames carrying MPLS-TP packets. Since MPLS-TP packets are MPLS packets, existing procedures ([RFC3032], [RFC5332]) for the encapsulation of MPLS packets over Ethernet apply. Because IP functionality is optional in an MPLS-TP network, IP-based protocols for Media Access Control (MAC) address learning, such as the Address Resolution Protocol (ARP) [RFC826] and IPv6 Neighbor Discovery [RFC4861], may not be available. This document specifies the options for the determination and selection of the next-hop Ethernet MAC address when MPLS-TP is used between nodes that do not have an IP data plane.

1.1. Terminology

Term	Definition
ARP	Address Resolution Protocol
G-ACh	Generic Associated Channel
LSP	Label Switched Path
LSR	Label Switching Router
MAC	Media Access Control
MPLS-TP	MPLS Transport Profile

Additional definitions and terminology can be found in [RFC5960] and [RFC5654].

1.2. 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 [RFC2119].

2. Point-to-Point Link Addressing

When two MPLS-TP nodes are connected by a point-to-point Ethernet link, the question arises as to what destination Ethernet Media Access Control (MAC) address should be specified in Ethernet frames transmitted to the peer node over the link. The problem of determining this address does not arise in IP/MPLS networks because

of the presence of the Ethernet Address Resolution Protocol (commonly referred to as the Address Resolution Protocol and shortened to ARP) [RFC826] or IPv6 Neighbor Discovery (ND) protocol [RFC4861], which allow the unicast MAC address of the peer device to be learned dynamically.

If existing mechanisms are available in an MPLS-TP network to determine the destination unicast MAC addresses of peer nodes, for example, if the network also happens to be an IP/MPLS network, or if the Link Layer Discovery Protocol (LLDP) [LLDP] is in use, these methods SHOULD be used. If ARP, ND, and LLDP are not available, and both peers implement the procedures in Section 4 of this document, then the GAP mechanism described in this memo SHOULD be used. The remainder of this section discusses alternative options that might be employed when none of the preceding mechanisms for learning MAC addresses are available.

One common approach is for each node to be statically configured with the MAC address of its peer. However, static MAC address configuration can present an administrative burden and lead to operational problems. For example, replacement of an Ethernet interface to resolve a hardware fault when this approach is used requires that the peer node be manually reconfigured with the new MAC address. This is especially problematic if the peer is operated by another provider.

Another approach that may be considered is to use the Ethernet broadcast address FF-FF-FF-FF-FF-FF as the destination MAC address in frames carrying MPLS-TP packets over a link that is known to be point-to-point. This may, however, lead to excessive frame distribution and processing at the Ethernet layer. Broadcast traffic may also be treated specially by some devices, and this may not be desirable for MPLS-TP data frames.

In view of the above considerations, this document recommends that, if a non-negotiated address is to be used, both nodes are configured to use as a destination MAC address an Ethernet multicast address reserved for MPLS-TP for use over point-to-point links. The address allocated for this purpose by the Internet Assigned Numbers Authority (IANA) is 01-00-5E-90-00-00. An MPLS-TP implementation MUST default to passing to the MPLS sub-system any MPLS packets received from a point-to-point Ethernet link with this destination MAC address.

The use of broadcast or multicast addressing for the purpose described in this section, i.e., as a placeholder for the unknown unicast MAC address of the destination, is applicable only when the attached Ethernet link is known to be point-to-point. If a link is not known to be point-to-point, these forms of broadcast or multicast

addressing MUST NOT be used. Thus, the implementation MUST provide a means for the operator to declare that a link is point-to-point if it supports these addressing modes. Moreover, the operator is cautioned that it is not always clear whether a given link is, or will remain, strictly point-to-point, particularly when the link is supplied by an external provider; point-to-point declarations therefore need to be used with care. Because of these caveats, it is RECOMMENDED that implementations support the procedures in Section 4 so that unicast addressing can be used.

3. Multipoint Link Addressing

When a multipoint Ethernet link serves as a section [RFC5960] for a point-to-multipoint MPLS-TP LSP, and multicast destination MAC addressing at the Ethernet layer is used for the LSP, the addressing and encapsulation procedures specified in [RFC5332] SHALL be used.

When a multipoint Ethernet link (that is, a link that is not known to be point-to-point) serves as a section for a point-to-point MPLS-TP LSP, unicast destination MAC addresses MUST be used for Ethernet frames carrying packets of the LSP. According to the discussion in the previous section, this implies the use of either static MAC address configuration or a protocol that enables peer MAC address discovery.


4. MAC Address Discovery via the G-ACh Advertisement Protocol

The G-ACh Advertisement Protocol (GAP) [RFC7212] provides a simple means of informing listeners on a link of the sender's capabilities and configuration. When used for this purpose on an Ethernet link, GAP messages are multicast to the address 01-00-5e-80-00-0d (see Section 7 of [RFC7212]). If these messages contain the unicast MAC address of the sender, then listeners can learn this address and use it in the future when transmitting frames containing MPLS-TP packets. Since the GAP does not rely on IP, this provides a means of unicast MAC discovery for MPLS-TP nodes without IP support.


This document defines a new GAP application "Ethernet Interface Parameters" (0x0001) to support the advertisement of Ethernet-specific parameters associated with the sending interface. The following Type-Length-Value (TLV) objects are defined for this application; the TLV format is as defined in [RFC7212]:

Source MAC Address (type = 0, length = 8): The Value of this object is an EUI-64 [EUI-64] unicast MAC address assigned to one of the interfaces of the sender that is connected to this data link. The IEEE-defined mapping from 48-bit MAC addresses to EUI-64 form is used.

Maximum Frame Size (MFS) (type = 1, length = 4): The Value of this object is a 32-bit unsigned integer encoded in network byte order that specifies the maximum frame size in octets of an Ethernet Frame that can be sent over the sending interface. Where MAC address learning occurs by some other means, this TLV group MAY be used to advertise only the MFS. If multiple advertisements are made for the same parameter, use of these advertisements is undefined.



Source MAC Address Object Format



MFS Object Format

Per [RFC7212], MAC address discovery information needs to be periodically retransmitted and is to be retained by a receiver based on the period of time indicated by the associated Lifetime field. To expedite the initialization of a link, it is RECOMMENDED that a node that has been reconfigured, rebooted, or is aware that it has been disconnected from its peer send a GAP Ethernet Interface Parameters message, and that it issue a GAP Request message for the Ethernet Interface Parameters of its peers, at the earliest opportunity.

When the MAC address in the received Source MAC Address TLV changes, the new MAC address MUST be used (see Section 5.2 of [RFC7212]).

If a minimum MFS is configured for a link and the MFS advertised by the peer is lower than that minimum, the operator MUST be notified of the MFS mismatch. Under these circumstances, the operator may choose to configure the LSR to shut down the link, thereby triggering a

fault and causing the end-to-end path to be repaired. Alternatively, the operator may choose to configure the LSR to leave the link up so that an OAM message can be used to verify the actual MFS.

A peer node could cease transmission of G-ACh advertisements, or cease to include a Source MAC Address TLV in advertisements, or cease to be connected, any of which will cause the TLV lifetime to expire. After the Source MAC Address TLV lifetime has expired, this MAC Address MUST NOT be used as the peer MAC address. The node MUST return to selecting MAC addresses as though no advertisements were received, using the method configured for this eventuality.

5. Manageability Considerations

The values sent and received by this protocol MUST be made accessible for inspection by network operators, and where local configuration is updated by received information, it MUST be clear why the configured value has been changed. If the received values change, the new values MUST be used and the change made visible to the network operators.

The Ethernet address and associated parameters advertised for an interface SHOULD be persistent across restarts. In the event of a system restart, any data that has been retained as a consequence of prior Ethernet Interface Parameters advertisements from GAP peers MUST be discarded; this prevents incorrect operation on the basis of stale data.

Where the link changes to a new type, i.e., from point-to-point to point-to-multipoint, the network operator SHOULD be informed. If the new link type is incompatible with the Ethernet addressing method in use, the system MUST take the action that is configured under those circumstances.

6. Security Considerations

The use of broadcast or multicast Ethernet destination MAC addresses for frames carrying MPLS-TP data packets can potentially result in such frames being distributed to devices other than the intended destination node or nodes when the Ethernet link is not point-to-point. The operator should take care to ensure that MPLS-TP nodes are aware of the Ethernet link type (point-to-point or multipoint). In the case of multipoint links, the operator should either ensure that no devices are attached to the link that are not authorized to receive the frames or take steps to mitigate the possibility of excessive frame distribution (for example, by configuring the Ethernet switch to appropriately restrict the delivery of multicast frames to authorized ports).

An attacker could disrupt communications by modifying the Source MAC Address or the MFS values; however, this is mitigated by the use of cryptographic authentication as described in [RFC7212], which also describes other considerations applicable to the GAP protocol. Visibility into the contents of either of the TLVs could provide information that is useful for an attacker. This is best addressed by physical security of the links.

7. IANA Considerations

7.1. Ethernet Multicast Address Allocation

IANA has allocated an Ethernet multicast address from the "IANA Multicast 48-bit MAC Addresses" address block in the "Ethernet Numbers" registry for use by MPLS-TP LSRs over point-to-point links as described in Section 2. The allocated address is 01-00-5E-90-00-00. IANA has updated the reference to point to the RFC number assigned to this document.

7.2. G-ACh Advertisement Protocol Allocation

IANA has allocated a new Application ID in the "G-ACh Advertisement Protocol Application Registry", as follows:

Application ID	Description	Reference
0x0001	Ethernet Interface Parameters	This RFC

7.3. Creation of Ethernet Interface Parameters Registry

IANA has created a new registry, "G-ACh Advertisement Protocol: Ethernet Interface Parameters" within the "Generic Associated Channel (G-ACh) Parameters" registry with fields and initial allocations as follows:

Type Name	Type ID	Reference
Source MAC Address	0	This RFC
Maximum Frame Size	1	This RFC

The range of the Type ID field is 0 - 255.

The allocation policy for this registry is IETF Review or IESG Approval.

8. Acknowledgements

We thank Adrian Farrel for his valuable review comments on this document.

9. References

9.1. Normative References

- [EUI-64] IEEE, "Guidelines for 64-bit Global Identifier (EUI-64) Registration Authority", March 1997, <<http://standards.ieee.org/regauth/oui/tutorials/EUI64.html>>.
- [LLDP] IEEE, "Station and Media Access Control Connectivity Discovery", IEEE 802.1AB, September 2009.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [RFC5332] Eckert, T., Rosen, E., Aggarwal, R., and Y. Rekhter, "MPLS Multicast Encapsulations", RFC 5332, August 2008.
- [RFC5654] Niven-Jenkins, B., Brungard, D., Betts, M., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, September 2009.
- [RFC5960] Frost, D., Bryant, S., and M. Bocci, "MPLS Transport Profile Data Plane Architecture", RFC 5960, August 2010.
- [RFC7212] Frost, D., Bryant, S., and M. Bocci, "MPLS Generic Associated Channel (G-ACh) Advertisement Protocol", RFC 7212, June 2014.

9.2. Informative References

- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC5921] Bocci, M., Bryant, S., Frost, D., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, July 2010.

[RFC826] Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", STD 37, RFC 826, November 1982.

Authors' Addresses

Dan Frost
Blue Sun

EEmail: frost@mm.st

Stewart Bryant
Cisco Systems

EEmail: stbryant@cisco.com

Matthew Bocci
Alcatel-Lucent

EEmail: matthew.bocci@alcatel-lucent.com

