

Network Working Group
Request for Comments: 5332
Category: Standards Track
Updates: 3032, 4023

T. Eckert
E. Rosen, Ed.
Cisco Systems, Inc.
R. Aggarwal
Y. Rekhter
Juniper Networks, Inc.
August 2008

MPLS Multicast Encapsulations

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

RFC 3032 established two data link layer codepoints for MPLS, used to distinguish whether the data link layer frame is carrying an MPLS unicast or an MPLS multicast packet. However, this usage was never deployed. This specification updates RFC 3032 by redefining the meaning of these two codepoints. Both codepoints can now be used to carry multicast packets. The second codepoint (formerly the "multicast codepoint") is now to be used only on multiaccess media, and it is to mean "the top label of the following label stack is an upstream-assigned label".

RFC 3032 does not specify the destination address to be placed in the "MAC DA" (Medium Access Layer Destination Address) field of an ethernet frame that carries an MPLS multicast packet. This document provides that specification.

This document updates RFC 3032 and RFC 4023.

Table of Contents

1. Introduction	2
2. Specification of Requirements	3
3. Upstream-Assigned vs. Downstream-Assigned	3
4. Ethernet Codepoints	6
5. PPP Protocol Field	6
6. GRE Protocol Type	6
7. IP Protocol Number	7
8. Ethernet MAC DA for Multicast MPLS	7
9. IANA Considerations	8
10. Security Considerations	8
11. Normative References	9

1. Introduction

RFC 3031 [RFC3031] defines the "Next Hop Label Forwarding Entry" (NHLFE). The NHLFE for a particular label maps the label into a next hop (among other things). When an MPLS packet is received, its top label is mapped to an NHLFE, and the packet is sent to the next hop specified by the NHLFE.

We define a particular MPLS label to be a "multicast label" in a particular context if the NHLFE to which it is mapped, in that context, specifies a set of next hops, with the semantics that the packet is to be replicated and a copy of the packet sent to each of the specified next hops. Note that this definition accommodates the case where the set of next hops contains a single member. What makes a label a multicast label in a particular context is the semantics attached to the set, i.e., the intention to replicate the packet and transmit to all members of the set if the set has more than one member.

RFC 3032 [RFC3032] established two data link layer codepoints for MPLS: one to indicate that the data link layer frame is carrying an MPLS unicast packet, and the other to indicate that the data link layer frame is carrying an MPLS multicast packet. The term "multicast packet" is not precisely defined in RFC 3032, though one may presume that the "multicast" codepoint is intended to identify the packet's top label as a multicast label. However, the multicast codepoint has never been deployed, and further development of the procedures for MPLS multicast have shown that, while there is a need for two codepoints, the use of the two codepoints is not properly captured by RFC 3032.

In particular, there is no need for the codepoint to indicate whether the top MPLS label is a multicast label. When the receiver of an MPLS packet looks up the top label, the NHLFE will specify whether or not the label is a multicast label.

This document updates RFC 3032 and RFC 4023 by re-specifying the use of the codepoints. The old use of the "multicast codepoint", as specified in those two RFCs, is hereby deprecated.

Note that an implementation that does MPLS multicast according to RFC 3032 and/or 4023 will be unable to interoperate with implementations that do MPLS multicast according to this document. There may be some deployed platforms that support the deprecated use of the codepoints, but those platforms do not support the control plane mechanisms to support MPLS multicast. The absence of the control plane will prevent a system that implements the deprecated use of codepoints from attempting to interoperate with a system that uses the codepoints as specified herein. (If an MPLS multicast control plane were to be implemented on a platform that only supports the deprecated codepoint, interoperability problems such as black holes and/or misrouting would arise. This does not seem like a potential problem in practice.)

While RFC 3032 allows an MPLS packet to be carried in an ethernet multicast frame, it fails to specify how the Medium Access Layer Destination Address (MAC DA) field is to be set in that case. This document provides that specification.

2. Specification of Requirements

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].

3. Upstream-Assigned vs. Downstream-Assigned

Suppose a labeled packet P is sent from Label Switching Router (LSR) R1 to LSR R2, where R1 puts label L on the packet's label stack, and R2 has to look up label L in order to determine the corresponding Forwarding Equivalence Class (FEC), call it F.

If the binding between L and F was made by R2 and advertised to R1, then the label binding is known as "downstream-assigned". RFC 3031 only discusses downstream-assigned label bindings.

If the binding between L and F was made by R1 and advertised to R2, then the label binding is known as "upstream-assigned".

If the binding between L and F was made by a third party, say R3, and then advertised to both R1 and R2, we also refer to the label binding as "upstream-assigned".

Upstream-assigned labels are not required to come from the same "label space" as downstream-assigned labels. See Section 3.14 of [RFC3031] and especially [RFC5331] for a discussion of the notion of "label space". The procedures for properly interpreting an upstream-assigned label are given in [RFC5331].

If Ru and Rd are LSP adjacencies, then they transmit an MPLS packet to each other through one of the following mechanisms:

1. by putting the MPLS packet in a data link layer frame and transmitting the frame,
2. by transmitting the MPLS packet through an MPLS tunnel, i.e., by pushing an additional label (or labels) onto the label stack, and then invoking mechanism 1, or
3. by transmitting the MPLS packet through an IP-based tunnel (e.g., via RFC 4023 [RFC4023]), and then invoking mechanisms 1 and/or 2.

In short, an MPLS packet is transmitted through a data link, through an MPLS tunnel, or through an IP tunnel. In any of those cases, when the packet emerges through the tunnel, the downstream LSR must know whether the label that now appears at the top of the label stack has an upstream-assigned label binding or a downstream-assigned label binding. For convenience, we will speak of a label with an upstream-assigned label binding as an "upstream-assigned label".

Under certain conditions, specified below, multicast labels MAY be upstream-assigned. The ability to use upstream-assigned labels is an OPTIONAL feature. Upstream-assigned labels MUST NOT be used unless it is known that the downstream LSR supports them. How this is known is outside the scope of this document.

This document makes no changes to the procedures regarding unicast labels.

We discuss three different types of data link or tunnel:

- Point-to-Point. A point-to-point data link or tunnel associates two systems, such that transmissions on that link or tunnel made by one are received by the other, and only by the other.

For a given direction of a given point-to-point data link or tunnel, the following MUST be the case: either every MPLS packet will carry an upstream-assigned label, or else every MPLS packet will carry a downstream-assigned label. The procedures for determining whether upstream-assigned or downstream-assigned labels are being used are outside the scope of this specification. However, in the absence of any other information, the use of downstream-assigned labels MUST be presumed by default.

- Point-to-Multipoint. A point-to-multipoint link or tunnel associates n systems, such that only one of them can transmit onto the link or tunnel, and the transmissions may be received by the other n-1 systems.

The top labels (before applying the data link or tunnel encapsulation) of all MPLS packets that are transmitted on a particular point-to-multipoint data link or tunnel MUST be of the same type; either all upstream-assigned or all downstream-assigned. This means that all the receivers on the MPLS or IP tunnel must know a priori whether upstream-assigned or downstream-assigned labels are being used in the tunnel. How this is known is outside the scope of this document.

- Multipoint-to-Multipoint. A multipoint-to-multipoint link or tunnel associates n systems, such that any of them can transmit on the link or tunnel, and the transmissions may be received by the other n-1 systems.

If MPLS packets are transmitted on a particular multipoint-to-multipoint link or tunnel, one of the following scenarios applies:

1. It is known (by methods outside the scope of this document) that the top label of every MPLS packet on the link or tunnel is downstream-assigned.
2. It is known (by methods outside the scope of this document) that the top label of every MPLS packet on the link or tunnel is upstream-assigned.
3. Some MPLS packets on the link may have upstream-assigned top labels while some may have downstream-assigned top labels.

If (and only if) the third scenario applies, the data link or tunnel encapsulation MUST provide a codepoint that specifies whether the top label of the encapsulated MPLS packet is upstream-assigned or downstream-assigned. If a particular type of data link or tunnel does not provide such a codepoint, then the third scenario MUST NOT be used.

The remainder of this document specifies procedures for setting the data link layer codepoints and address fields.

4. Ethernet Codepoints

Ethernet is an example of a multipoint-to-multipoint data link.

Ethertype 0x8847 is used whenever a unicast ethernet frame carries an MPLS packet.

Ethertype 0x8847 is also used whenever a multicast ethernet frame carries an MPLS packet, EXCEPT for the case where the top label of the MPLS packet has been upstream-assigned.

Ethertype 0x8848, formerly known as the "MPLS multicast codepoint", is to be used only when an MPLS packet whose top label is upstream-assigned is carried in a multicast ethernet frame.

5. PPP Protocol Field

PPP is an example of a point-to-point data link. When a PPP frame is carrying an MPLS packet, the PPP Protocol field is always set to 0x0281.

6. GRE Protocol Type

RFC 4023 is modified as described below.

If the IP destination address of the Generic Routing Encapsulation (GRE) is a unicast IP address, then the ethertype value 0x8847 MUST be used in all cases for the MPLS-in-GRE encapsulation.

If the IP destination address of the GRE encapsulation is a multicast IP address, then:

- the ethertype value 0x8847 MUST be used when the top label of the encapsulated MPLS packet is downstream-assigned,
- the ethertype value 0x8848 MUST be used when the top label of the encapsulated MPLS packet is upstream-assigned.

Through procedures that are outside the scope of this specification, it may be known that if the destination address of a GRE packet is a multicast IP address, then the top label of the GRE payload is upstream-assigned. In such a case, the occurrence of the 8847 codepoint in a GRE packet with a multicast destination IP address MUST be considered an error, and the packet MUST be discarded.

7. IP Protocol Number

RFC 4023 is modified as follows: the IPv4 Protocol Number field or the IPv6 Next Header field is always set to 137, whether or not the encapsulated MPLS packet is an MPLS multicast packet.

If the IP destination address of the IP encapsulation is an IP multicast address, the IP tunnel may be considered to be a point-to-multipoint tunnel or a multipoint-to-multipoint tunnel. In either case, either all encapsulated MPLS packets in the particular tunnel have a downstream-assigned label at the top of the stack, or all encapsulated MPLS packets in that tunnel have an upstream-assigned label at the top of the stack. The means by which this is determined for a particular tunnel is outside the scope of this specification.

8. Ethernet MAC DA for Multicast MPLS

When an LSR transmits a multicast MPLS packet in a multicast ethernet frame, it MUST set the MAC Destination Address to the value 01-00-5e-8v-wx-yz, where vwxyz is a 20-bit (five-nibble) value set as follows:

1. vwxyz MAY be set to 0,
2. vwxyz MAY be set to the value of one of the MPLS labels on the packet's label stack.

Which of these procedures is the default procedure in any particular LSR is implementation-dependent. However, LSRs using the two different procedures MUST interoperate. That is, an LSR MUST NOT filter packets for which vwxyz has been set to zero, and it MUST NOT indiscriminately filter all packets for which vwxyz has not been set to zero.

If an LSR follows the procedure of setting vwxyz to the value of one of the MPLS labels on the packet's label stack, and if that label stack contains two or more labels, then by default, vwxyz MUST be set to the value of the second MPLS label on the packet's label stack. By "the second label", we mean the label that is in the label stack entry that immediately follows the topmost label stack entry. The LSR MAY, if configured to do so, allow a label other than the second

to be used for this purpose. If the MPLS packet has only one label, the value of that label will be used instead of the value of the (non-existent) second label.

It is expected that the LSR will follow the procedures of [RFC5331], pushing on two labels, with the topmost label being a "context label" that is the same for all MPLS packets being transmitted by the LSR onto the ethernet, but with the second label being different for different LSPs. Thus, if the MAC DA value is a function of the second label, more of the LSP-specific information about the packet appears in the MAC DA field. This can be used to filter multicast packets with "unexpected" non-zero values of vwxyz. Further discussion of such filtering or its uses is outside the scope of this document.

The use of ethernet and/or IP broadcast addresses (as distinguished from multicast addresses) does not fall within the scope of this specification.

9. IANA Considerations

IANA already owns the set of ethernet multicast addresses in the range 01-00-5e-00-00-00 to 01-00-5e-ff-ff-ff. Addresses in the range 01-00-5e-00-00-00 to 01-00-5e-7f-ff-ff are already reserved for use when an ethernet multicast frame carries an IP multicast packet.

IANA has reserved ethernet addresses in the range 01-00-5e-80-00-00 to 01-00-5e-8f-ff-ff for use when an ethernet multicast frame carries an MPLS multicast packet. Addresses in this range are valid when used with ethertype 8847 or 8848.

As this document modifies the usage of ethertypes 8847 and 8848, IANA has changed the description of these ethertypes as follows. Ethertype 8847 is defined as "MPLS", as defined in RFC 3032 and in this document. Ethertype 8848 is defined as "MPLS with upstream-assigned label", as defined in this document.

10. Security Considerations

The security considerations of RFC 3032 and RFC 4023 apply.

Malicious changing of the codepoint may result in loss or misrouting of packets. However, altering the codepoint without also altering the label does not result in a predictable effect.

Malicious alteration of the MAC DA on an ethernet can result in packets being received by a third party, rather than by the intended recipient.

11. 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.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, March 2005.
- [RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", RFC 5331, August 2008.

Authors' Addresses

Toerless Eckert
Cisco Systems, Inc.
170 Tasman Drive
San Jose, CA, 95134

EEmail: eckert@cisco.com

Eric C. Rosen
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719

EEmail: erosen@cisco.com

Rahul Aggarwal
Juniper Networks
1194 North Mathilda Ave.
Sunnyvale, CA 94089

EEmail: rahul@juniper.net

Yakov Rekhter
Juniper Networks
1194 North Mathilda Ave.
Sunnyvale, CA 94089

EEmail: yakov@juniper.net

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.