

MPLS Upstream Label Assignment and Context-Specific Label Space

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 3031 limits the MPLS architecture to downstream-assigned MPLS labels. This document introduces the notion of upstream-assigned MPLS labels. It describes the procedures for upstream MPLS label assignment and introduces the concept of a "Context-Specific Label Space".

Table of Contents

1. Introduction	2
2. Specification of Requirements	2
3. Context-Specific Label Space	2
4. Upstream Label Assignment	3
4.1. Upstream-Assigned and Downstream-Assigned Labels	4
5. Assigning Upstream-Assigned Labels	5
6. Distributing Upstream-Assigned Labels	5
7. Upstream Neighbor Label Space	6
8. Context Label on LANs	9
9. Usage of Upstream-Assigned Labels	10
10. Security Considerations	10
11. Acknowledgements	11
12. References	11
12.1. Normative References	11
12.2. Informative References	11

1. Introduction

RFC 3031 [RFC3031] limits the MPLS architecture to downstream-assigned MPLS labels. To quote from RFC 3031:

"In the MPLS architecture, the decision to bind a particular label L to a particular Forwarding Equivalence Class (FEC) F is made by the Label Switching Router (LSR) which is DOWNSTREAM with respect to that binding. The downstream LSR then informs the upstream LSR of the binding. Thus labels are "downstream-assigned", and label bindings are distributed in the "downstream to upstream" direction."

This document introduces the notion of upstream-assigned MPLS labels to the MPLS architecture. The procedures for upstream assignment of MPLS labels are described.

RFC 3031 describes per-platform and per-interface label space. This document generalizes the latter to a "Context-Specific Label Space" and describes a "Neighbor Label Space" as an example of this. Upstream-assigned labels are always looked up in a context-specific label space.

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. Context-Specific Label Space

RFC 3031 describes per-platform and per-interface label spaces. This document introduces the more general concept of a "Context-Specific Label Space". An LSR may maintain one or more context-specific label spaces. In general, labels MUST be looked up in the per-platform label space unless something about the context determines that a label be looked up in a particular context-specific label space.

One example of a context-specific label space is the per-interface label space discussed in RFC 3031. When an MPLS packet is received over a particular interface, the top label of the packet may need to be looked up in the receiving interface's per-interface label space. In this case, the receiving interface is the context of the packet. Whether MPLS packets received over a particular interface need to have their top labels looked up in a per-interface label space depends on some characteristic or configuration of the interface.

Per-interface label space [RFC3031] is an example of a context-specific label space used for downstream-assigned labels. Context-specific label spaces can also be used for upstream-assigned labels, as described below.

When MPLS labels are upstream-assigned, the context of an MPLS label L is provided by the LSR that assigns the label and binds the label to a FEC F for a Label Switched Path (LSP) LSP1. The LSR that assigns the label distributes the binding and context to an LSR Lr that then receives MPLS packets on LSP1 with label L. When Lr receives an MPLS packet on LSP1, it MUST be able to determine the context of this packet.

An example of such a context is a tunnel over which MPLS packets on LSP1 may be received. In this case, the top label of the MPLS packet, after tunnel decapsulation, is looked up in a label space that is specific to the root of the tunnel. This does imply that Lr be able to determine the tunnel over which the packet was received. Therefore, if the tunnel is an MPLS tunnel, penultimate-hop-popping (PHP) MUST be disabled for the tunnel.

Another example of such a context is the neighbor from which MPLS packets on LSP1 may be received. In this case, the top label of the MPLS packet, transmitted by the neighbor on LSP1, is looked up in a "Neighbor-Specific Label Space".

The above two examples are further described in Section 7.

There may be other sorts of contexts as well. For instance, we define the notion of an MPLS label being used to establish a context, i.e., identify a label space. A "context label" is one that identifies a label table in which the label immediately below the context label should be looked up. A context label carried as an outermost label over a particular multi-access subnet/tunnel MUST be unique within the scope of that subnet/tunnel.

4. Upstream Label Assignment

When two MPLS LSRs are adjacent in an MPLS Label Switched Path (LSP), one of them can be termed an "upstream LSR" and the other a "downstream LSR" [RFC3031]. Consider two LSRs, Ru and Rd, that have agreed to bind Label L to a FEC F for packets sent from Ru to Rd. Then, with respect to this binding, Ru is the "upstream LSR", and Rd is the "downstream LSR".

If the binding between L and F was made by Rd and advertised to Ru, 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 Ru and advertised to Rd, 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 Ru and Rd, we also refer to the label binding as "upstream-assigned".

An important observation about upstream-assigned labels is the following. When an upstream-assigned label L is at the top of the label stack, it must be looked up by an LSR that is not the LSR that assigned and distributed the label binding for L. Therefore, an upstream-assigned label MUST always be looked up in a context-specific label space, as described in Section 7.

We do not require any coordination between the upstream label assignments and the downstream label assignments; a particular label value may be upstream-assigned to one FEC and downstream-assigned to a different FEC.

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.

One use case of upstream-assigned labels is MPLS multicast, and an example of this is provided in Section 9.

4.1. Upstream-Assigned and Downstream-Assigned Labels

It is possible that some LSRs on an LSP for FEC F distribute downstream-assigned label bindings for FEC F, while other LSRs distribute upstream-assigned label bindings. It is possible for an LSR to distribute a downstream-assigned label binding for FEC F to its upstream adjacent LSR AND distribute an upstream-assigned label binding for FEC F to its downstream adjacent LSR. When two LSRs, Ru and Rd, are adjacent on an LSP for FEC F (with Ru being the upstream neighbor and Rd the downstream neighbor), either Ru distributes an upstream-assigned label binding for F to Rd, or else Rd distributes a downstream-assigned label binding to Ru, but NOT both. Whether upstream-assigned or downstream-assigned labels are to be used for a particular FEC depends on the application using the LSP.

Any application that requires the use of upstream-assigned labels MUST specify that explicitly, or else it is to be assumed that downstream-assigned labels are used. An application on an LSR uses a label distribution protocol to indicate to its peer LSRs whether a particular label binding distributed by the LSR uses upstream-assigned or downstream-assigned label. Details of such procedures are outside the scope of this document. In some cases, the decision

as to which is used for a particular application may be made by a configuration option.

5. Assigning Upstream-Assigned Labels

The only requirement on an upstream LSR assigning upstream-assigned labels is that an upstream-assigned label must be unambiguous in the context-specific label space in which the downstream LSR will look it up. An upstream LSR that is the headend of multiple tunnels SHOULD, by default, assign the upstream-assigned labels, for all the LSPs carried over these tunnels, from a single label space, which is common to all those tunnels. Further, an upstream LSR that is the head of multiple tunnels SHOULD use the same IP address as the head identifier of these tunnels, provided that the head identifier of these tunnels includes an IP address. The LSR could assign the same label value to both a downstream-assigned and an upstream-assigned label. The downstream LSR always looks up upstream-assigned MPLS labels in a context-specific label space as described in Section 7.

An entry for the upstream-assigned labels is not created in the Incoming Label Map (ILM) [RFC3031] at the upstream LSR as these labels are not incoming labels. Instead, an upstream label is an outgoing label, with respect to the upstream LSR, for MPLS packets transmitted on the MPLS LSP in which the upstream LSR is adjacent to the downstream LSR. Hence, an upstream label is part of a Next Hop Label Forwarding Entry (NHLFE) at the upstream LSR.

When R_u advertises a binding of label L for FEC F to R_d , it creates a NHLFE entry corresponding to L . This NHLFE entry results in imposing the label L on the MPLS label stack of the packet forwarded using the NHLFE entry. If R_u is a transit router on the LSP for FEC F , it binds the ILM for the LSP to this NHLFE. If R_u is an ingress router on the LSP for FEC F , it binds the FEC to the NHLFE entry.

6. Distributing Upstream-Assigned Labels

Upstream-assigned label bindings 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.

MPLS upstream label assignment requires a label distribution protocol to distribute the binding from the upstream LSR to the downstream LSR. Considerations that pertain to a label distribution protocol that are described in [RFC3031] apply.

The distribution of the upstream-assigned labels is similar to either the ordered LSP control or independent LSP control of the downstream-assigned labels. In the former case, an LSR distributes an upstream-

assigned label binding for a FEC F if it is either (a) the ingress LSR for FEC F, or (b) if it has already received an upstream label binding for that FEC from its adjacent upstream LSR for FEC F, or (c) if it has received a request for a downstream label binding from its upstream adjacent LSR. In the latter case, each LSR, upon noting that it recognizes a particular FEC, makes an independent decision to bind an upstream-assigned label to that FEC and to distribute that binding to its label distribution peers.

7. Upstream Neighbor Label Space

If the top label of an MPLS packet being processed by LSR Rd is upstream-assigned, the label is looked up in a context-specific label space, not in a per-platform label space.

Rd uses a context-specific label space that it maintains for Ru to "reserve" MPLS labels assigned by Ru. Hence, if Ru distributes an upstream-assigned label binding L for FEC F to Rd, then Rd reserves L in the separate ILM for Ru's context-specific label space. This is the ILM that Rd uses to look up an MPLS label that is upstream-assigned by Ru. This label may be the top label on the label stack of a packet received from Ru or it may be exposed as the top label on the label stack, as a result of Rd popping one or more labels off the label stack, from such a packet.

This implies that Rd MUST be able to determine whether the top label of an MPLS packet being processed is upstream-assigned and, if yes, the "context" of this packet. How this determination is made depends on the mechanism that is used by Ru to transmit the MPLS packet with an upstream-assigned top label L to Rd.

If Ru transmits this packet by encapsulating it in an IP or MPLS tunnel, then the fact that L is upstream-assigned is determined by Rd by the tunnel on which the packet is received. Whether a given tunnel can be used for transmitting MPLS packets with either downstream-assigned or upstream-assigned MPLS labels, or both, depends on the tunnel type and is described in [RFC5332]. When Rd receives MPLS packets with a top label L on such a tunnel, it determines the "context" of this packet based on the tunnel on which the packet is received. There must be a mechanism for Ru to inform Rd that a particular tunnel from Ru to Rd will be used by Ru for transmitting MPLS packets with upstream-assigned MPLS labels. Such a mechanism will be provided by the label distribution protocol between Ru and Rd and will likely require extensions to existing label distribution protocols. The description of such a mechanism is outside the scope of this document.

Rd maintains an "Upstream Neighbor Label Space" for upstream-assigned labels, assigned by Ru. When Ru transmits MPLS packets the top label of which is upstream-assigned over IP or MPLS tunnels, then Rd MUST be able to determine the root of these IP/MPLS tunnels. Rd MUST then use a separate label space for each unique root.

The root is identified by the head-end IP address of the tunnel. If the same upstream router, Ru, uses different head-end IP addresses for different tunnels, then the downstream router, Rd, MUST maintain a different Upstream Neighbor Label Space for each such head-end IP address.

Consider the following conditions:

- 1) Ru is the "root" of two tunnels, call them A and B.
- 2) IP address X is an IP address of Ru.
- 3) The signaling protocol used to set up tunnel A identified A's root node as IP address X.
- 4) The signaling protocol used to set up tunnel B identified B's root node as IP address X.
- 5) Packets sent through tunnels A and B may be carrying upstream-assigned labels.
- 6) Ru is the LSR that assigned the upstream-assigned labels mentioned in condition 5.

If and only if these conditions hold, then Ru MUST use the same label space when upstream-assigning labels for packets that travel through tunnel A that it uses when upstream-assigning labels for packets that travel through tunnel B.

Suppose that Rd is a node that belongs to tunnels A and B, but is not the root node of either tunnel. Then Rd may assume that the same upstream-assigned label space is used on both tunnels IF AND ONLY IF the signaling protocol used to set up tunnel A identified the root node as IP address X and the signaling protocol used to set up tunnel B identified the root node as the same IP address X.

In addition, the protocol that is used for distributing the upstream-assigned label to be used over a particular tunnel MUST identify the "assigner" using the same IP address that is used by the protocol that sets up the tunnel to identify the root node of the tunnel. Implementors must take note of this, even if the tunnel setup

protocol is different from the protocol that is used for distributing the upstream-assigned label to be used over the tunnel.

The precise set of procedures for identifying the IP address of the root of the tunnel depend, of course, on the protocol used to set up the tunnel. For Point-to-Point (P2P) tunnels, the intention is that the headend of the tunnel is the "root". For Point-to-Multipoint (P2MP) or Multipoint-to-Multipoint (MP2MP) tunnels, one can always identify one node as being the "root" of the tunnel.

Some tunnels may be set up by configuration, rather than by signaling. In these cases, the IP address of the root of the tunnel must be configured.

Some tunnels may not even require configuration, e.g., a Generic Routing Encapsulation (GRE) tunnel can be "created" just by encapsulating packets and transmitting them. In such a case, the IP address of the root is considered to be the IP source address of the encapsulated packets.

If the tunnel on which Rd receives MPLS packets with a top label L is an MPLS tunnel, then Rd determines a) that L is upstream-assigned and b) the context for L, from the labels above L in the label stack. Note that one or more of these labels may also be upstream-assigned labels.

If the tunnel on which Rd receives MPLS packets with a top label L is an IP/GRE tunnel, then Rd determines a) that L is upstream-assigned [RFC5332] and b) the context for L, from the source address in the IP header.

When Ru and Rd are adjacent to each other on a multi-access data link media, if Ru would transmit the packet, with top label L, by encapsulating it in a data link frame, then whether L is upstream-assigned or downstream assigned can be determined by Rd, as described in [RFC5332]. This is possible because if L is upstream-assigned, then [RFC5332] uses a different ether type in the data link frame. However, this is not sufficient for Rd to determine the context of this packet. In order for Rd to determine the context of this packet, Ru encapsulates the packet in a one-hop MPLS tunnel. This tunnel uses an MPLS context label that is assigned by Ru. Section 8 describes how the context label is assigned. Rd maintains a separate "Upstream Neighbor Label Space" for Ru. The "context" of this packet, i.e., Ru's upstream neighbor label space, in which L was reserved, is determined by Rd from the top context label and the interface on which the packet is received. The ether type in the data link frame is set to indicate that the top label is upstream-assigned. The second label in the stack is L.

8. Context Label on LANs

For a labeled packet with an ether type of "upstream label assignment", the top label is used as the context. The context label value is assigned by the upstream LSR and advertised to the downstream LSRs. Mechanisms for advertising the context label will be provided by the label distribution protocol between the upstream and downstream LSRs. The description of such a mechanism is outside the scope of this document.

The context label assigned by an LSR for use on a particular LAN interface MUST be unique across all the context labels assigned by other LSRs for use on the same LAN. When a labeled packet is received from the LAN, the context label MUST be looked up in the context of the LAN interface on which the packet is received.

This document provides two methods that an LSR can use to choose a context label to advertise on a particular LAN.

The first method requires that each LSR be provisioned with a 20-bit context label for each LAN interface on which a context label is required. It is then left to the provisioning system to make sure that an assigned context label is unique across the corresponding LAN.

The second method allows the context labels to be auto-generated, but is only applicable if each LSR on the LAN has an IPv4 address as its primary IP address for the corresponding LAN interface. (If the LAN contains LSRs that have only IPv6 addresses for the LAN interface, then the first method is used.)

Suppose that each LAN interface is configured with a primary IPv4 address that is unique on that LAN. The host part of the IPv4 address, identified by the network mask, is unique. If the IPv4 network mask is greater than 12 bits, it is possible to map the remaining 20 bits into a unique context label value. This enables the LSRs on the LAN to automatically generate a unique context label. To ensure that auto-generated context label values do not fall into the reserved label space range [RFC3032], the value of the host part of the IPv4 address is offset with 0x10, if this value is not greater than 0xFFFFEF. Values of the host part of the IPv4 address greater than 0xFFFFEF are not allowed to be used as context labels.

Consider LSR Rm (downstream) connected to Ru1 (upstream) on a LAN interface and to Ru2 (upstream) on a different LAN interface. Rm could receive a context label value derived from the LAN interface from Ru1 and from Ru2. It is possible that the context label values used by Ru1 and Ru2 are the same. This would occur if the LAN

interfaces of both Ru1 and Ru2 are configured with a primary IPv4 address where the lowest 20 bits are equal. However, this does not create any ambiguity, as it has already been stated that the context label MUST be looked up in the context of the LAN interface on which the packet is received.

9. Usage of Upstream-Assigned Labels

A typical use case of upstream-assigned labels is for MPLS multicast and is described here for illustration. This use case arises when an upstream LSR Ru is adjacent to several downstream LSRs <Rd1..Rdn> in an LSP, LSP1 AND Ru is connected to <Rd1..Rdn> via a multi-access media or tunnel, AND Ru wants to transmit a single copy of an MPLS packet on the LSP to <Rd1..Rdn>. In the case of a tunnel, Ru can distribute an upstream-assigned label L that is bound to the FEC for LSP1, to <Rd1..Rdn> and transmit an MPLS packet, the top label of which is L, on the tunnel. In the case of a multi-access media, Ru can distribute an upstream-assigned label L that is bound to the FEC for LSP1, to <Rd1..Rdn> and transmit an MPLS packet, the top label of which is the context label that identifies Ru, and the label immediately below is L, on the multi-access media. Each of <Rd1..Rdn> will then interpret this MPLS packet in the context of Ru and forward it appropriately. This implies that <Rd1..Rdn> MUST all be able to support an Upstream Neighbor Label Space for Ru and Ru MUST be able to determine this. The mechanisms for determining this are specific to the application that is using upstream-assigned labels and is outside the scope of this document.

10. Security Considerations

The security considerations that apply to upstream-assigned labels and context labels are no different in kind than those that apply to downstream-assigned labels.

Note that procedures for distributing upstream-assigned labels and/or context labels are not within the scope of this document. Therefore, the security considerations that may apply to such procedures are not considered here.

Section 8 of this document describes a procedure that enables an LSR to automatically generate a unique context label for a LAN. This procedure assumes that the IP addresses of all the LSR interfaces on the LAN will be unique in their low-order 20 bits. If two LSRs whose IP addresses have the same low-order 20 bits are placed on the LAN, other LSRs are likely to misroute packets transmitted to the LAN by either of the two LSRs in question.

More detailed discussion of security issues that are relevant in the context of MPLS and GMPLS, including security threats, related defensive techniques, and the mechanisms for detection and reporting, are discussed in "Security Framework for MPLS and GMPLS Networks [MPLS-SEC].

11. Acknowledgements

Thanks to IJsbrand Wijnands's contribution, specifically for the text on which Section 8 is based.

12. References

12.1. Normative References

- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5332] Eckert, T., Rosen, E., Aggarwal, R., and Y. Rekhter, "MPLS Multicast Encapsulations", RFC 5332, August 2008.

12.2. Informative References

- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [MPLS-SEC] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", Work in Progress, July 2008.

Authors' Addresses

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

EMail: rahul@juniper.net

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

EMail: yakov@juniper.net

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

EMail: erosen@cisco.com

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.

