

Network Working Group

Internet-Draft

Updates: [3032](#), [3038](#), [3209](#), [3811](#), [4182](#), [4928](#), [5331](#),
[5586](#), [5921](#), [5960](#), [6391](#), [6478](#), [6790](#) (if approved)

Intended status: Standards Track

Expires: August 14, 2014

K. Kompella

Juniper Networks

L. Andersson

Huawei

A. Farrel

Juniper Networks

February 14, 2014

Allocating and Retiring Special Purpose MPLS Labels draft-ietf-mpls-special-purpose-labels-05

Abstract

Some MPLS labels have been allocated for specific purposes. A block of labels (0-15) has been set aside to this end, and are commonly called "reserved labels". They will be called "special purpose labels" in this document.

As there are only 16 of these special purpose labels, caution is needed in the allocation of new special purpose labels, yet at the same time allow forward progress when one is called for.

This memo defines new procedures to follow in the allocation and retirement of special purpose labels, as well as a method to extend the special purpose label space. Finally, this memo renames the IANA registry for these labels to "Special Purpose MPLS Label Values", and creates a new one called the "Extended Special Purpose MPLS Label Values" registry.

This document updates a number of previous RFCs that used the term "reserved label".

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 <http://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 August 14, 2014.

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.

Table of Contents

1.	Introduction	3
1.1.	Conventions used	3
2.	Questions	4
3.	Answers	5
3.1.	Extended Special Purpose MPLS Label Values	6
3.2.	Process for Retiring Special Purpose Labels	7
4.	Updated RFCs	8
5.	IANA Considerations	9
6.	Security Considerations	10
7.	Acknowledgments	10
8.	References	11
8.1.	Normative References	11
8.2.	Informational References	12
	Authors' Addresses	13

1. Introduction

The specification of the Label Stack Encoding for Multi-Protocol Label Switching (MPLS) [[RFC3032](#)] defined four special purpose label values (0 to 3), and set aside values 4 through 15 for future use. These labels have special significance in both the control and the data plane. Since then, three further values have been allocated (values 7, 13, and 14 in [[RFC6790](#)], [[RFC5586](#)] and [[RFC3429](#)], respectively), leaving nine unassigned values from the original space of sixteen.

While the allocation of three out of the remaining twelve special purpose label values in the space of about 12 years is not in itself a cause for concern, the scarcity of special purpose labels is. Furthermore, many of the special purpose labels require special processing by forwarding hardware, changes to which are often expensive, and sometimes impossible. Thus, documenting a newly allocated special purpose label value is important.

This memo outlines some of the issues in allocating and retiring special purpose label values, and defines mechanisms to address these. This memo also extends the space of special purpose labels.

1.1. Conventions used

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](#)].

Two new acronyms are introduced:

XL The Extension Label that indicates that an extended special purpose label follows.

ESPL An Extended Special Purpose Label. A Special Purpose Label that is placed in the label stack after the Extension Label. The combination of XL and ESPL might be regarded as a new form of "compound label" comprising more than one consecutive entry in the label stack.

2. Questions

In re-appraising MPLS special purpose labels, the following questions come to mind:

1. What allocation policies should be applied by IANA for the allocation of special purpose labels? Should Early Allocation [[RFC7120](#)] be allowed? Should there be labels for Experimental Use or Private Use [[RFC5226](#)]?
2. What documentation is required for special purpose labels allocated henceforth?
3. Should a special purpose label ever be retired? What criteria are relevant here? Can a retired special purpose label ever be re-allocated for a different purpose? What procedures and time frames are appropriate?
4. The special purpose label value of 3 (the "Implicit Null Label", [[RFC3032](#)]) is only used in signaling, never in the data plane. Could it (and should it) be used in the data plane? If so, how and for what purpose?
5. What is a feasible mechanism to extend the space of special purpose labels, should this become necessary?
6. Should extended special purpose labels be used for load balancing?

3. Answers

This section provides answers to the questions posed in the previous section.

1.

- A. Allocation of special purpose MPLS labels is via "Standards Action".
- B. The IANA registry will be renamed "Special Purpose MPLS Labels".
- C. Early allocation may be allowed on a case-by-case basis.
- D. The current space of 16 special purpose labels is too small for setting aside values for experimental or private use. However, the extended special purpose labels registry created by this document has enough space, and this document defines a range for experimental use.

- 2. A Standards Track RFC must accompany a request for allocation of Standards Action special purpose labels, as per [[RFC5226](#)].
- 3. The retirement of a special purpose MPLS label value must follow a strict and well-documented process. This is necessary since we must avoid orphaning the use of this label value in existing deployments. This process is detailed in [Section 3.2](#).
- 4. For now, the use of the "implicit null label" (label 3) in the data plane will not be allowed. If this decision is revisited later, an accompanying Standards Track RFC that details the use of the label, a discussion of possible sources of confusion between signaling and data plane, and mitigation thereof shall be required.
- 5. A special purpose label (the "extension" label, XL, label 15) is set aside for the purpose of extending the space of special purpose labels. Further details are described in [Section 3.1](#).
- 6. [[RFC6790](#)] says that special purpose labels MUST NOT be used for load balancing. The same logic applies to extended special purpose labels (ESPLs). Thus, this document specifies that ESPLs MUST NOT be used for load balancing. It is noted that existing implementations would violate this, as they do not recognize XL as anything other than a single Special Purpose Label and will not expect an ESPL to follow. The consequence is that if ESPLs are used in some packets of a flow, these packets may be

delivered on different paths and so could be re-ordered. However, it is important to specify the correct behavior for future implementations, hence the use of "MUST NOT".

A further question that needed to be settled in this regard was whether a "regular" special purpose label retains its meaning if it follows the XL. This answer to this question is provided in [Section 3.1](#).

[3.1](#). Extended Special Purpose MPLS Label Values

The XL MUST be followed by another label L (and thus MUST have the bottom-of-stack bit clear). L MUST be interpreted as an ESPL and interpreted as defined in a new registry created by this document (see [Section 5](#)). Whether or not L has the bottom-of-stack bit set depends on whether other labels follow L. The XL only assigns special meaning to L. A label after L (if any) is parsed as usual, and thus may be a regular label or a special purpose label; if the latter, it may be the XL, and thus followed by another ESPL.

The label value 15 is set aside as the XL as shown in [Section 5](#).

Values 0-6 and 8-15 of the Extended Special Purpose Label registry are set aside as reserved; these MUST NOT appear in the data plane. If an LSR encounters such ESPL values it MUST treat the packet as malformed and discard it per [\[RFC3031\]](#).

Label 7 (when received) retains its meaning as ELI whether a regular special purpose label or an ESPL; this simplifies a transit LSR's task of looking for entropy labels since it may just look for label 7 and need not verify that the previous label in the stack is not the XL 15. However, an LSR wishing to insert an entropy label SHOULD insert label 7 as a regular special purpose label, not as an ESPL.

[3.1.1](#). Forwarding Packets with Extended Special Purpose Labels

If an LSR encounters the XL at the top of stack and it doesn't understand extension labels, it SHOULD drop the packet as specified for the handling of any unknown label according to [\[RFC3031\]](#). If an LSR encounters an ESPL at the top of stack (after the XL) and does not understand the ESPL, it SHOULD drop the packet, again following the procedures for unknown labels as set out in [\[RFC3031\]](#). In either case, the LSR MAY log the event, but such logging MUST be rate-limited.

An LSR SHOULD NOT make forwarding decisions on labels not at the top of stack. For load balancing decisions, see Answer 6 of [Section 3](#).

3.1.2. Choosing a New Special Purpose Label

When allocating a new Special Purpose Label, protocol designers should consider whether they could use an Extended Special Purpose Label. Doing so would help to preserve the scarce resources of "normal" Special Purpose Labels for use in cases where minimizing the label stack size is particularly important.

3.2. Process for Retiring Special Purpose Labels

While the following process is defined for the sake of completeness, note that retiring special purpose labels is difficult. It is recommended that this process be used sparingly.

- a. A label value that has been assigned from the "Special Purpose MPLS Label Values" may be deprecated by IETF consensus with review by the MPLS working group (or designated experts if the working group or a successor does not exist). An RFC with at least Informational status is required.

The RFC will direct the IANA to mark the label value as "deprecated" in the registry, but will not release it at this stage.

Deprecating means that no further specifications using the deprecated value will be documented.

At the same time this is an indication to vendors not to include the deprecated value in new implementations, and to operators to avoid including it in new deployments.

- b. 12 months after the RFC deprecating the label value is published, an IETF-wide survey may be conducted to determine if the deprecated label value is still in use. If the survey indicates that the deprecated label value is in use, the survey may be repeated after a further 6 months.
- c. 24 months after the RFC that deprecated the label value was published and if the survey indicates that deprecated label value is not in use, publication may be requested of an IETF Standards Track Internet-Draft that retires the deprecated the label value. This document will request IANA to release the label value for future use and assignment.

4. Updated RFCs

The following RFCs contain references to the term "reserved labels":

- o [[RFC3032](#)] ("MPLS Label Stack Encoding"),
- o [[RFC3038](#)] ("VCID Notification for LDP")
- o [[RFC3209](#)] ("Extensions to RSVP for LSP Tunnels")
- o [[RFC3811](#)] ("MPLS TC MIB")
- o [[RFC4182](#)] ("Removing a Restriction on the use of MPLS")
- o [[RFC4928](#)] ("Avoiding ECMP Treatment in MPLS Networks")
- o [[RFC5331](#)], [[RFC5586](#)] ("G-ACh and GAL")
- o [[RFC5921](#)] ("MPLS Transport Profile Framework")
- o [[RFC5960](#)] ("MPLS-TP Data Plane Architecture")
- o [[RFC6391](#)] ("FAT-PW")
- o [[RFC6478](#)] ("Pseudowire Status for Static Pseudowires")
- o [[RFC6790](#)] ("MPLS Entropy Labels").

All such references should be read as "special purpose labels".

5. IANA Considerations

This document requests IANA to make the following changes and additions to its registration of MPLS Labels.

1. Change the name of the "Multiprotocol Label Switching Architecture (MPLS) Label Values" registry to the "Special Purpose MPLS Label Values".
2. Change the allocations policy for the "Special Purpose MPLS Label Values" registry to Standards Action.
3. Assign label 15 from the "Special Purpose MPLS Label Values" registry, naming it the "extension label", and citing this document as the reference.
4. Create a new registry called the "Extended Special Purpose MPLS Label Values" registry. The ranges and allocation policies for this registry are as follows in Table 1 (using terminology from [\[RFC5226\]](#)). Early allocation following the policy defined in [\[RFC7120\]](#) is allowed only for those values assigned by Standards Action.

Range	Allocation Policy
0 - 6, 8 - 15	Reserved. Not to be allocated. MUST NOT appear in the data plane.
7	Allocated; meaning is ELI [RFC6790]
16 - 239	Standards Action
240 - 255	Experimental
256 - 1048575	Reserved. Not to be allocated without a new Standards Track RFC to define an allocation policy.

Table 1

6. Security Considerations

This document does not make a large change to the operation of the MPLS data plane and security considerations are largely unchanged from those specified in the MPLS architecture [[RFC3031](#)] and in the MPLS and GMPLS Security Framework [[RFC5920](#)].

However, it should be noted that increasing the label stack can cause packet fragmentation and may also make packets unprocessable by some implementations. This document provides a protocol-legal way to increase the label stack through the insertion of additional {XL,ESPL} pairs at a greater rate than insertion of single "rogue" labels. This might provide a way to attack some nodes in a network that can only process label stacks of a certain size without violating the protocol rules.

This document also describes events that may cause an LSR to issue event logs at a per-packet rate. It is critically important that implementations rate-limit such logs.

7. Acknowledgments

Thanks to Pablo Frank and Lizhong Jin for useful discussions.
Thanks to Curtis Villamizar, Mach Chen, Alia Atlas, Eric Rosen, and Maria Napierala for useful comments.

8. References

8.1. 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.
- [RFC3038] Nagami, K., Katsube, Y., Demizu, N., Esaki, H., and P. Doolan, "VCID Notification over ATM link for LDP", [RFC 3038](#), January 2001.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.
- [RFC3811] Nadeau, T. and J. Cucchiara, "Definitions of Textual Conventions (TCs) for Multiprotocol Label Switching (MPLS) Management", [RFC 3811](#), June 2004.
- [RFC4182] Rosen, E., "Removing a Restriction on the use of MPLS Explicit NULL", [RFC 4182](#), September 2005.
- [RFC4928] Swallow, G., Bryant, S., and L. Andersson, "Avoiding Equal Cost Multipath Treatment in MPLS Networks", [BCP 128](#), [RFC 4928](#), June 2007.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.
- [RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", [RFC 5331](#), August 2008.
- [RFC5960] Frost, D., Bryant, S., and M. Bocci, "MPLS Transport Profile Data Plane Architecture", [RFC 5960](#), August 2010.
- [RFC6391] Bryant, S., Filsfils, C., Drafz, U., Kompella, V., Regan, J., and S. Amante, "Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network", [RFC 6391](#), November 2011.

- [RFC6478] Martini, L., Swallow, G., Heron, G., and M. Bocci, "Pseudowire Status for Static Pseudowires", [RFC 6478](#), May 2012.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", [RFC 6790](#), November 2012.
- [RFC7120] Cotton, M., "Early IANA Allocation of Standards Track Code Points", [BCP 100](#), [RFC 7120](#), January 2014.

8.2. Informational References

- [RFC3429] Ohta, H., "Assignment of the 'OAM Alert Label' for Multiprotocol Label Switching Architecture (MPLS) Operation and Maintenance (OAM) Functions", [RFC 3429](#), November 2002.
- [RFC5586] Bocci, M., Vigoureux, M., and S. Bryant, "MPLS Generic Associated Channel", [RFC 5586](#), June 2009.
- [RFC5920] Fang, L., "Security Framework for MPLS and GMPLS Networks", [RFC 5920](#), July 2010.
- [RFC5921] Bocci, M., Bryant, S., Frost, D., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", [RFC 5921](#), July 2010.

Authors' Addresses

Kireeti Kompella
Juniper Networks
1194 N. Mathilda Ave
Sunnyvale, CA 94089
US

Email: kireeti.kompella@gmail.com

Loa Andersson
Huawei

Email: loa@mail01.huawei.com

Adrian Farrel
Juniper Networks

Email: adrian@olddog.co.uk