

Network Working Group
Internet-Draft
Updates: [5586](#) (if approved)
Intended status: Standards Track
Expires: March 27, 2016

A. Vainshtein
ECI Telecom
L. Andersson
Huawei Technologies Co., Ltd
A. Farrel
Juniper Networks
September 24, 2015

Handling the TC and TTL fields in a Label Stack Entry that Contains the
Generic Associated Channel Label
[draft-vainshtein-mpls-gal-tc-ttl-handling-03](#)

Abstract

This document clarifies handling of the Traffic Class (TC) and Time-to-Live (TTL) fields of a Label Stack Entry that contains the Generic Associated Channel (G-ACh) Label (GAL). These clarifications are intended to aid interoperability of implementations.

Original handling was defined in [RFC 5586](#), and this document updates that RFC.

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 March 27, 2016.

Copyright Notice

Copyright (c) 2015 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	2
2.	Terminology	4
2.1.	Requirements Language	4
2.2.	Abbreviations	4
3.	New Procedures	5
3.1.	New Procedures for Handling the TC Field in an LSE That Contains the GAL	5
3.2.	New Procedures for Handling the TTL Field in an LSE Containing GAL	5
3.3.	Scope of the new Procedures	5
4.	IANA Considerations	5
5.	Security Considerations	6
6.	Acknowledgments	6
7.	References	6
7.1.	Normative References	6
7.2.	Informative References	7
	Authors' Addresses	7

[1.](#) Introduction

[RFC5586] introduced an alert mechanism for the Generic Associated Channel (G-ACh) that uses a Generic Associated Channel Label (GAL). In particular, [RFC5586] allocated one of the values from the special purpose label space to be the GAL, specified that the Label Stack Entry (LSE) containing GAL must be always at the bottom of the label stack in the case of MPLS transport profile (MPLS-TP) Label Switched Paths (LSPs), and that G-ACh packets must not be forwarded based on the GAL.

Per [RFC3032] each LSE contains, in addition to the label value and bottom-of-stack (BoS) flag, two additional fields:

- o Traffic Class (TC) field - 3 bits (renamed from Experimental (EXP) field [RFC5462]). [RFC5586] defined that the handling of this field in an LSE that contains the GAL is as specified and referenced in [RFC 5462](#).
- o Time-to-Live (TTL) field - 8 bits. [RFC5586] defined that the handling of this field in an LSE that contains the GAL is in

accordance with [\[RFC3443\]](#). In addition, it defined (in [Section 4.2.1.1 of \[RFC3443\]](#)) that:

- * The TTL field of the GAL LSE MUST be set to at least 1. The exact value of the TTL is application specific.

The above-mentioned references to [RFC 5462](#) and [RFC 3443](#) are not very useful for the implementers because they do not define specific processing actions for these fields. That is, even when they give guidance on how a sending implementation should set a field, they don't explain how a receiving implementation should be "liberal in what it accepts". For example, [\[RFC5462\]](#) says only that the use of TC field for Quality of Service (QoS) and Explicit Congestion Notification (ECN) "is intended to be flexible". On the other hand, while [\[RFC3443\]](#) is very detailed with regard to processing of the TTL field, it mainly deals with issues that are irrelevant for an LSE that contains the GAL because [RFC 5586](#) explicitly prohibits forwarding labeled packets based on the GAL.

Implementations of [\[RFC5586\]](#) have encountered interoperability problems in their interpretation of these two fields when present in an LSE that contains the GAL.

When this LSE becomes the top entry in the label stack (because the previous label has been popped) some receiving implementations have attempted to interpret the TC and TTL fields. In particular, packets with an LSE that contains the GAL with TTL set to one can be trapped to the generic (slow) MPLS exception handler with appropriate rate limiting before the GAL is noticed (which would otherwise result in trapping the packet to a fast OAM handler). The resultant poor performance in handling TTL of one could be seen as a poor implementation choice, but it does not violate [RFC 5586](#).

Similarly, the setting of the TC field in the top LSE impacts how a packet is forwarded. When the top LSE contains the GAL, the packet will not be forwarded, so the value of the TC field should not be relevant (or, at most, should impact only on internal prioritization such as within a software implementation). However, some implementations might action the TC field before determining that the GAL is present and so might be subject to poor or unexpected behavior dependent on the setting of the field.

This document clarifies the rules for setting and processing the TC and TTL fields in the LSE that contains the GAL in an unambiguous way without referring to any other documents. This will improve interoperability with implementations that are potentially flawed in their handling of these fields. It updates [\[RFC5586\]](#) in that regard.

In summary, the objective of this document is not to change the conformance requirements of existing implementations, but to enhance the likelihood of new implementations interoperating with each other and with implementations already in the field. In particular, to improve the interoperability results where a new implementation interacts with a legacy implementation that may be sender or receiver of an LSE that contains the GAL and that may be any of:

- o built with poor design considerations
- o implemented with a misunderstanding of [RFC 5586](#)
- o over-zealous in its policing of [RFC 5586](#) behavior.

This document does not require existing implementations to change their behavior.

[2.](#) Terminology

[2.1.](#) 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](#)].

[2.2.](#) Abbreviations

BoS: Bottom of Stack

G-ACh: Generic Associated Channel

GAL: Generic Associated Channel Label

LER: Label Edge Router

LSE: Label Stack Entry

LSP: Label Switching Path

LSR: Label Switching Router

PW: Pseudowire

TC: Traffic Class field (formerly named EXP)

TTL: Time-to-Live

3. New Procedures

3.1. New Procedures for Handling the TC Field in an LSE That Contains the GAL

Setting the value of the TC field in an LSE that contains the GAL is done by the LER that originates the G-ACh packet and is a matter of local policy for that LER. It is RECOMMENDED that implementations set the TC field of an LSE that contains the GAL to all zero (0b000), but they MAY vary this according to local policy.

The LER that inspects an LSE that contains the GAL MUST ignore the value of the TC field.

3.2. New Procedures for Handling the TTL Field in an LSE Containing GAL

Setting the value of the TTL in an LSE that contains the GAL is done by the LER that originates the G-ACh packet and is a matter of local policy for that LER. The LER that originates the G-ACh packet MUST NOT set this value to 0 and SHOULD NOT set this value to 1: this will avoid possible misinterpretation by the LER that inspects an LSE that contains the GAL if that LER does not comply with this document. It is RECOMMENDED that implementations set the TTL of an LSE that contains the GAL to 255, but they MAY vary this according to local policy.

An LER that examines an LSE that contains the GAL MUST ignore the value of the TTL field.

3.3. Scope of the new Procedures

[RFC5586] disallowed the use of the GAL in PWS, but that limitation was relaxed in [[RFC6423](#)].

The new procedures defined in this document for handling the TC field and the TTL field in an LSE that contains the GAL apply equally to all possible uses of the GAL including the so-called "Section G-ACh" where the GAL is the only label in the label stack, and the use of the GAL in LSPs and PWS.

4. IANA Considerations

This document makes no requests for IANA action.

5. Security Considerations

This document makes a minor update to the processing for MPLS packets containing the GAL and does not change any of the security fundamentals of MPLS. For a discussion of security considerations relating to MPLS, please refer to [[RFC5920](#)].

Note that the rules set out in this document specify that a receiver must ignore the values in the two MPLS LSE fields that are discussed. As such, this clarification removes a potential (and minor) attack vector where those fields could be malignly set and might cause incorrect action by the receiver.

6. Acknowledgments

The authors would like to thank Nadav Baadany whose question triggered this work, and Jeff Haas, Jie Dong, Mach Chen, Huub van Helvoort, Carlos Pignataro, Curtis Villamizar, George Swallow, Rolf Winter, and Martin Vigoureux for their comments on this document.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", [RFC 3032](#), DOI 10.17487/RFC3032, January 2001, <<http://www.rfc-editor.org/info/rfc3032>>.
- [RFC3443] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", [RFC 3443](#), DOI 10.17487/RFC3443, January 2003, <<http://www.rfc-editor.org/info/rfc3443>>.
- [RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", [RFC 5462](#), DOI 10.17487/RFC5462, February 2009, <<http://www.rfc-editor.org/info/rfc5462>>.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", [RFC 5586](#), DOI 10.17487/RFC5586, June 2009, <<http://www.rfc-editor.org/info/rfc5586>>.

[RFC6423] Li, H., Martini, L., He, J., and F. Huang, "Using the Generic Associated Channel Label for Pseudowire in the MPLS Transport Profile (MPLS-TP)", [RFC 6423](#), DOI 10.17487/RFC6423, November 2011, <<http://www.rfc-editor.org/info/rfc6423>>.

7.2. Informative References

[RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", [RFC 5920](#), DOI 10.17487/RFC5920, July 2010, <<http://www.rfc-editor.org/info/rfc5920>>.

Authors' Addresses

Alexander Vainshtein
ECI Telecom
Petah Tikva
Israel

Email: alexander.vainshtein@ecitele.com

Loa Andersson
Huawei Technologies Co., Ltd
Stockholm
Sweden

Email: loa@mail01.huawei.com

Adrian Farrel
Juniper Networks

Email: adrian@olddog.co.uk

