

ROLL
Internet-Draft
Intended status: Standards Track
Expires: May 10, 2013

P. Thubert, Ed.
Cisco
November 6, 2012

Use of the IPv6 Flow Label within an LLN
draft-thubert-roll-flow-label-02

Abstract

This document present how the Flow Label can be used inside a LLN as a replacement to the RPL option and provides rules for the root to set and reset the Flow Label when forwarding between the inside of RPL domain and the larger Internet, in both direction. This new operation aims at saving an IPv6 in IPv6 encapsulation within the RPL domain that is required with the RPL option for all packets that reach outside of the RPL domain.

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 May 10, 2013.

Copyright Notice

Copyright (c) 2012 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

Internet-Draft Use of the IPv6 Flow Label within an LLN November 2012

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Terminology	5
3.	Flow Label Format Within the RPL Domain	5
4.	Root Operation	6
4.1.	Incoming Packets	6
4.2.	Outgoing Packets	7
5.	RPL node Operation	7
6.	Security Considerations	7
7.	IANA Considerations	7
8.	Acknowledgments	7
9.	References	7
9.1.	Normative References	7
9.2.	Informative References	8
	Author's Address	8

Internet-Draft Use of the IPv6 Flow Label within an LLN November 2012

1. Introduction

In some Low Power and Lossy Network (LLN) applications such as control systems [[RFC5673](#)], a packet loss is usually acceptable but jitter and latency must be strictly controlled as they can play a critical role in the interpretation of the measured information. Sensory systems are often distributed, and the control information can in fact be originated from multiple sources and aggregated. As a result, it can be a requirement for related measurements from multiple sources to be treated as a single flow following a same path over the Internet in order to experience similar jitter and latency. The traditional tuple of source, destination and ports might then not be the proper indication to isolate a meaningful flow.

In a typical LLN application, the bulk of the traffic consists of small chunks of data (in the order few bytes to a few tens of bytes) at a time. In the industrial case, a typical frequency is 4Hz but it can be a lot slower than that for, say, environmental monitoring. The granularity of traffic from a single source is too small to make a lot of sense in load balancing application.

In such cases, related packets from multiple sources should not be load-balanced along their path in the Internet; load-balancing can be discouraged by tagging those packets with a same Flow Label in the IPv6 [[RFC2460](#)] header. This can be achieved if the Flow Label in packets outgoing a RPL domain are set by the root of the RPL structure as opposed to the actual source. It derives that the Flow Label could be reused inside the RPL domain.

In a LLN, each transmitted bit represents energy and every saving counts dearly. Considering that the value for which the Flow Label is used in the IPv6 Flow Label Specification [[RFC6437](#)] is to serve load balancing in the core, it is unlikely that LLN devices will consume energy to generate and then transmit a Flow Label to serve interests in some other place. On the other hand, it makes sense to recommend the computation of a stateless Flow Label at the root of

the LLN towards the Internet.

Reciprocally, [RFC6437] requires that once set, a non-zero flow label value is left unchanged. The value for that setting is consumed once the packet has traversed the core and reaches the LLN. Then again, there is little value but a high cost for the LLN in spending 20 bits to transport a Flow Label from the Internet over the constrained network to the destination node. It results that the MUST in [RFC6437] should be alleviated for packets coming from the outside on the LLN, and that it should be acceptable that the compression over the LLN erases the original flow label. It should also be acceptable that the Flow Label field is reused in the LLN as proposed in this

Thubert

Expires May 10, 2013

[Page 3]

Internet-Draft Use of the IPv6 Flow Label within an LLN November 2012

draft.

The Routing Protocol for Low Power and Lossy Networks (RPL) [[RFC6550](#)] specification defines a generic Distance Vector protocol that is adapted to a variety of LLNs. RPL forms Destination Oriented Directed Acyclic Graphs (DODAGs) which root often acts as the Border Router to connect the RPL domain to the Internet. The root is responsible to select the RPL Instance that is used to forward a packet coming from the Internet into the RPL domain.

A classical RPL implementation will use the RPL Option for Carrying RPL Information in Data-Plane Datagrams [[RFC6553](#)] to tag a packet with the Instance ID and other information that RPL requires for its operation within the RPL domain. Sadly, the Option must be placed in a Hop-by-Hop header that must be added to or removed from packets that cross the border of the RPL domain. For reasons such as the capability to send ICMP errors, back, this operation involves an extra 6in6 encapsulation within the RPL domain that is detrimental to the LLN operation, in particular with regards to bandwidth and battery constraints. The extra encapsulation may cause a containing frame to grow above maximum frame size, leading to Layer 2 or 6LoWPAN [[RFC4944](#)] fragmentation, which in turn cause even more energy spending and issues discussed in the LLN Fragment Forwarding and Recovery [[I-D.thubert-roll-forwarding-frags](#)].






Figure 1: 6in6 Encapsulation within the LLN

Additionally, Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks [[RFC6282](#)] and its variants for other types of

LLNs do not provide an efficient compression for the RPL option so the cost in current implementations can not be alleviated in any fashion. So even for packets that are confined within the RPL domain and do not need the 6in6 encapsulation, the use of the flow label instead of the RPL option is a valuable saving.

All the packets that are leaving a DODAG of a RPL domain towards the Internet will transit via a same root. The root is an ideal place to set the IPv6 Flow Label to a same value across multiple sources of a same flow when that operation is needed, ensuring compliance with the rules defined by the IPv6 Flow Label Specification [[RFC6437](#)] within the Internet. At the same time, the root segregates the Internet and the RPL domain, allowing to reuse the Flow Label within the RPL domain.

It can be noted that [[RFC6282](#)] provides an efficient header compression for packets that do have the Flow Label set in the IPv6 header. It results that the same information as transported in the RPL option itself represents actually less bits in the air when the Flow Label is used instead. This document specifies how the Flow Label can be reused within the RPL domain as a replacement to the RPL option. The use of the Flow Label within a RPL domain is an instance

of the stateful scenarios as discussed in [[RFC6437](#)] where the states include the rank of a node and the RPLInstanceID that identifies the routing topology.

2. Terminology

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](#)].

The Terminology used in this document is consistent with and incorporates that described in 'Terminology in Low power And Lossy Networks' [[I-D.ietf-roll-terminology](#)] and [[RFC6550](#)].

3. Flow Label Format Within the RPL Domain

[[RFC6550](#)] [section 11.2](#) specifies the fields that are to be placed into the packets for the purpose of Instance Identification, as well as Loop Avoidance and Detection. Those fields include an 'O', and 'R' and an 'F' bits, the 8-bit RPLInstanceID, and the 16-bit SenderRank. SenderRank is the result of the DAGRank operation on the rank of the sender, where the DAGRank operation is defined in [section 3.5.1](#) as:

Thubert

Expires May 10, 2013

[Page 5]

Internet-Draft Use of the IPv6 Flow Label within an LLN November 2012

$$\text{DAGRank}(\text{rank}) = \text{floor}(\text{rank} / \text{MinHopRankIncrease})$$

If MinHopRankIncrease is set to a multiple of 256, it appears that the most significant 8 bits of the SenderRank will be all zeroes and could be omitted. In that case, the Flow Label MAY be used as a replacement to the [[RFC6553](#)] RPL option. To achieve this, the SenderRank is expressed with 8 least significant bits, and the information carried within the Flow Label in a packet is constructed follows:

0	1	2						
0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5 6 7 8 9	0 1 2 3						
+--+								
O R F SenderRank RPLInstanceID								

+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

Figure 1: The RPL Flow Label

The first (leftmost) bit of the Flow Label is reserved and should be set to zero.

[4.](#) Root Operation

[RFC6437] [section 3](#) intentionally does not consider flow label values in which any of the bits have semantic significance. However, the present specification assigns semantics to various bits in the flow label, destroying within the edge network that is the RPL domain a property of belonging to a statistically uniform distribution that is desirable in the rest of the Internet. This property MUST be restored by the root for outgoing packets.

It can be noted that the rationale for the statistically uniform distribution does not necessarily bring a lot of value within the RPL domain. In a specific use case where it would, that value must be compared with that of the battery savings in order to decide which technique the deployment will use to transport the RPL information.

[4.1.](#) Incoming Packets

When routing a packet towards the RPL domain, the root applies a policy to determine whether the Flow Label is to be used to carry the RPL information. If so, the root MUST reset the Flow Label and then it MUST set all the fields in the Flow Label as prescribed by [\[RFC6553\]](#) using the format specified in Figure 1. In particular, the root selects the Instance that will be used to forward the packet

within the RPL domain.

[4.2.](#) Outgoing Packets

When routing a packet outside the RPL domain, the root applies a policy to determine whether the Flow Label was used to carry the RPL information. If so, the root MUST reset the Flow Label. The root SHOULD recompute a Flow Label following the rules prescribed by

[RFC6553]. In particular, the root MAY ignore the source address but it SHOULD use the RPLInstanceID for the computation.

5. RPL node Operation

Depending on the policy in place, the source of a packet will decide whether to use this specification to transport the RPL information in the IPv6 packets. If it does, the source in the LLN SHOULD set the Flow Label to zero and MUST NOT expect that the flow label will be conserved end-to-end".

6. Security Considerations

The process of using the Flow Label as opposed to the RPL option does not appear to create any opening for new threat compared to [RFC6553].

7. IANA Considerations

No IANA action is required for this specification.

8. Acknowledgments

The author wishes to thank Brian Carpenter for his in-depth review and constructive approach to the problem and its resolution.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.

Datagrams over IEEE 802.15.4-Based Networks", [RFC 6282](#), September 2011.

- [RFC6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", [RFC 6550](#), March 2012.
- [RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", [RFC 6553](#), March 2012.

[9.2.](#) Informative References

- [I-D.ietf-roll-terminology]
Vasseur, J., "Terminology in Low power And Lossy Networks", [draft-ietf-roll-terminology-06](#) (work in progress), September 2011.
- [I-D.thubert-roll-forwarding-frags]
Thubert, P. and J. Hui, "LLN Fragment Forwarding and Recovery", [draft-thubert-roll-forwarding-frags-00](#) (work in progress), March 2012.
- [RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", [RFC 4944](#), September 2007.
- [RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", [RFC 5673](#), October 2009.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", [RFC 6437](#), November 2011.

Author's Address

Pascal Thubert (editor)
Cisco Systems
Village d'Entreprises Green Side
400, Avenue de Roumanille
Batiment T3
Biot - Sophia Antipolis 06410
FRANCE

Phone: +33 4 97 23 26 34
Email: pthubert@cisco.com

Thubert

Expires May 10, 2013

[Page 9]