

6MAN
Internet-Draft
Intended status: Standards Track
Expires: February 24, 2015

P. Thubert, Ed.
Cisco
August 25, 2014

The IPv6 Flow Label within a LLN domain
draft-thubert-6man-flow-label-for-rpl-05

Abstract

This document presents how the Flow Label can be used inside a LLN domain such as a RPL domain or an ISA100.11a D-subnet, and provides updated rules for a domain Border Router to set and reset the Flow Label when forwarding between inside the domain and the larger Internet in both direction. Rules for routers inside the domain are also provided.

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 February 24, 2015.

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	2
2. Terminology	3

3. Requirements for LLN Flows 3

4. On Compatibility With Existing Standards 4

5. Updated Rules 5

6. Security Considerations 6

7. IANA Considerations 6

8. Acknowledgements 6

9. References 7

 9.1. Normative References 7

 9.2. Informative References 7

Author's Address 8

1. Introduction

The design of Lowpower Lossy Networks (LLNs) is generally focussed on saving energy, which is typically the most constrained resource of all. Other classical constraints, such as memory capacity, frame size, as well as the duty cycling of the LLN devices, derive from that primary concern.

In isolated devices, energy is typically available from batteries that are expected to last for years, or scavenged from the environment in very limited quantities. Any protocol that is intended for use in LLNs must be designed with the primary concern of saving energy as a strict requirement.

The IEEE802.15.4 [IEEE802154] was designed to offer the Physical (PHY) and Medium Access Control (MAC) layers for low-cost, low-speed, low-power Wireless Personal Area Networks (WPANs), which are a wireless form of LLNs.

With the traditional IEEE802.15.4 PHY, frames are limited to 127 octets. In order to adapt IPv6 [RFC2460] over IEEE802.15.4, 6LoWPAN [RFC4944] introduced a fragmentation mechanism under IP, which in turn causes even more energy spending and other issues as discussed in LLN Fragment Forwarding and Recovery [I-D.thubert-6lo-forwarding-fragments].

The IEEE802.15.4e Task Group further defined the TimeSlotted Channel Hopping [I-D.ietf-6tisch-tsch] (TSCH) mode of operation as an update to the MAC specification in order to address Time Sensitive applications.

The 6TISCH architecture [I-D.ietf-6tisch-architecture] specifies the operation of IPv6 over IEEE802.15.4e TSCH networks attached and synchronized by backbone routers. 6TiSCH was created to simplify the adoption of IETF technology by other Standard Defining Organizations (SDOs), in particular in the Industrial Automation space, which already relies on variations of IEEE802.15.4e TSCH for Wireless Sensor Networking.

The ISA100.11a [ISA100.11a] specification provides an example of such an industrial WSN standard, using a precursor to IEEE802.15.4e over the classical IEEE802.14.5 PHY. In that case, after security is applied, roughly 80 octets are available per frame for IP and Payload. In order to 1) avoid fragmentation and 2) conserve energy, the ISA100 WG in charge of that specification did scrutinize the use of every bit in the frame and rejected any perceived waste.

The challenge to obtain the adoption of IPv6 in the original standard was thus to save all possible bits in the frames, including the UDP checksum which was an interesting discussion on its own. This work was actually one of the roots for the 6LoWPAN Header Compression [RFC6282] work, which goes down to the individual bits to save space in the frames for actual data, and allowed ISA100.11a to adopt IPv6.

ISA100.11a (now IEC62734) uses IPv6 over UDP, and conforms to a number of other IETF RFCs including the IPv6 Flow Label Specification [RFC3697] that was the reference at the time the standard was elaborated, but fails to conform to the newer IPv6 Flow Label Specification [RFC6437] that obsoleted it.

The bone of contention is the use of the Flow Label as an index called a contract ID, and the capability for the Backbone Router, that is the Border Router of a ISA100.11a WSN (also called a D-subnet), to modify the Flow Label. There is work at ROLL that indicates that RPL nodes may benefit from similar abilities to also transport flow-related information in the Flow Label.

This document adds an exception to the rules in [RFC6437], for application within a well-defined LLN domain, whereby the Border Routers would be in a position to ensure that from an external viewpoint, the domain complies to the new Flow Label specification even though the internal use of the Flow Label does not.

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].

This document uses Terminology defined in Terminology in Low power And Lossy Networks [RFC7102], as well as [RFC6550] and [RFC6553].

3. Requirements for LLN Flows

In Industrial Automation and Control Systems (IACS) [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. In such cases, related packets from multiple sources should not be load-balanced

along their path in the Internet.

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. 4Hz is a typical loop frequency in Process Control, though it can be a lot slower than that in, say, environmental monitoring. The granularity of traffic from a single source is too small to make a lot of sense in load balancing application.

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 so as to experience similar jitter and latency. The traditional tuple of source, destination and ports might then not be the proper indication to isolate a consistent flow. On the other hand, the flow integrity can be preserved in a simple manner if the setting of the Flow Label in the IPv6 header of packets outgoing a LLN domain, is centralized to the Border Router, such as the root of a RPL DODAG structure, or an ISA100.11a Backbone Router, as opposed to distributed across the actual sources.

Considering that the goal for setting the Flow Label as prescribed in the IPv6 Flow Label Specification [RFC6437] is to improve load balancing in the core of the Internet, it is unlikely that LLN devices will consume energy to generate and then transmit a Flow Label to serve outside interests and the Flow Label is generally left to zero so as to be elided in the 6LoWPAN [RFC6282] compression. So in a general manner the interests of the core are better served if the RPL roots systematically rewrite the flow label rather than if they never do.

For packets coming into the RPL domain from the Internet, the value for setting the Flow Label as prescribed in [RFC6437] 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, that was set by a peer or a router in the Internet, over the constrained network to a destination node that has no use of it.

On a PHY layer with super-short frames such as IEEE802.15.4, compliance with those rules will simply not happen, and the rules will become an bone of contention for IPv6 adoption at a time where great progress is happening towards that goal, as illustrated by the activity at 6lo on multiple LLN Link-layers.

4. On Compatibility With Existing Standards

All the packets from all the nodes in a same DODAG that are leaving a RPL domain towards the Internet will transit via a same RPL root. The RPL root segregates the Internet and the RPL domain, which enables the capability to reuse the Flow Label within the RPL domain. The ISA100.11a Backbone Router plays a similar role and interfaces an ISA100.11a WSN D-subnet with a larger IPv6 network.

This specification enables the operation of resetting or reusing the IPv6 Flow Label at the border of a LLN domain. This is a deviation from the IPv6 Flow Label Specification [RFC6437], in that the LLN border router is neither the source nor the first hop router that sets the final Flow Label for use outside the LLN domain.

But if we consider the whole RPL domain as a large virtual host from the standpoint of the rest of the Internet, the interests that lead to [RFC6437], and in particular load balancing in the core of the Internet, are probably better served if the root guarantees that the Flow Label is set in a compliant fashion than if we rely on each individual sensor that may not use it at all, or use it slightly differently such as done in ISA100.11a.

Additionally, LLN flows can be compound flows aggregating information from multiple sources. The Border Router is an ideal place to rewrite the Flow Label to a same value for a same flow across multiple sources, ensuring compliance with the rules defined by [RFC6437] for use outside of the RPL domain and in particular in the core of the Internet.

This document specifies how the Flow Label can be reused within a LLN domain such as a RPL domain and an ISA100.11a D-subnet, in which a Border Router delineates the limit of the domain and may rewrite the Flow Label on all packets. In a RPL domain, it will become acceptable to use the Flow Label as replacement to the RPL option, though whether that operation gets standardized is left to be discussed. That use of the Flow Label within a RPL domain would be an instance of the stateful scenarios as discussed in [RFC6437] where the flow state in the node is indexed by the RPLInstanceID that identifies the routing topology. ISA100.11a would be another instance where the 16bit Contract ID in the Flow Label identifies a state in a node that is specific to a particular flow.

5. Updated Rules

This specification applies to a constrained LLN domain that forms a stub and is connected to the Internet by and only by its Border Routers. In the case of a RPL domain, the RPL root is such a bottleneck for all the traffic between the Internet and the Destination-Oriented Directed Acyclic Graph (DODAG) that it serves. This specification also covers other LLN domains with the same properties of having strict constraints in energy and/or frame size, such as an ISA100.11a [ISA100.11a] Industrial Wireless Sensor Network, but does not generalize to any arbitrary domain. This updates the IPv6 Flow Label Specification [RFC6437], which does not allow any specific rule in any particular domain, and updates it only in the context of constrained LLN domains.

In that context, a LLN domain Border Router MAY rewrite the Flow Label of all packets entering or leaving the RPL domain in both directions, from and towards the Internet, regardless of its original setting. For the limited context of a constrained LLN domain, this updates the IPv6 Flow Label Specification [RFC6437] which stipulates that once it is set, the Flow Label is left unchanged; but the RFC also indicates a violation to the rule can be accepted for compelling reasons related to security. This specification adds that energy-saving is another compelling reason for a violation to the aforementioned rule, though applicable only inside a constrained LLN.

In particular, the Border Router of a LLN domain MAY set the Flow Label of IPv6 packets that exit the LLN domain. It SHOULD do it if the LLN domain operations do not conform [RFC6437], and if it does modify the Flow Label, then it MUST do it in a manner that conforms [RFC6437] from the perspective of a Node outside the LLN.

It results that a Node in a constrained LLN domain MUST NOT assume that the setting of the Flow Label will be preserved end-to-end, and that an intermediate router inside a constrained LLN MAY alter a non-zero Flow Label between the source in the LLN and the LLN Border Router. This does not modify the expectations on end Nodes but extends the updated rules from [RFC6437] to arbitrary routers in the LLN.

For instance, a RPL root MAY reset the Flow Label of IPv6 packets entering the RPL domain to zero for an optimal Header Compression by 6LoWPAN [RFC6282]. A RPL root MAY also reuse the Flow Label towards the LLN for other purposes, such as to carry the RPL Information [RFC6553]. An ISA100.11s Backbone Router MAY reuse the Flow Label to carry local flow information, such as the Contract ID specified in ISA100.11a [ISA100.11a].

6. Security Considerations

Because the flow label is not protected by IPSec, it is expected that Layer-2 security is deployed in the LLN where its specification is applied. This is the actual best practice in LLNs, which serves in particular to avoid forwarding of untrusted packets over the constrained network.

The specification insists that the LLN Node should not expect that the Flow Label is conserved end-to-end and rather reduces the risk of misinterpretation in case of a rewrite by a router in the middle.

7. IANA Considerations

No IANA action is required for this specification.

8. Acknowledgements

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

9. References

9.1. Normative References

- [IEEE802154]
IEEE standard for Information Technology, "IEEE std. 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks", June 2011.
- [ISA100.11a]
ISA/ANSI, "Wireless Systems for Industrial Automation: Process Control and Related Applications - ISA100.11a-2011 - IEC 62734", 2011, <<http://www.isa.org/Community/SP100WirelessSystemsforAutomation>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2460] Deering, S.E. and R.M. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC3697] Rajahalme, J., Conta, A., Carpenter, B. and S. Deering, "IPv6 Flow Label Specification", RFC 3697, March 2004.
- [RFC6282] Hui, J. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, September 2011.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S. and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, November 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.
- [RFC6552] Thubert, P., "Objective Function Zero for the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6552, 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-6tisch-architecture]
Thubert, P., Watteyne, T. and R. Assimiti, "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4e", Internet-Draft draft-ietf-6tisch-architecture-01, February 2014.

[I-D.ietf-6tisch-tsch]

Watteyne, T., Palattella, M. and L. Grieco, "Using IEEE802.15.4e TSCH in an LLN context: Overview, Problem Statement and Goals", Internet-Draft draft-ietf-6tisch-tsch-00, November 2013.

[I-D.thubert-6lo-forwarding-fragments]

Thubert, P. and J. Hui, "LLN Fragment Forwarding and Recovery", Internet-Draft draft-thubert-6lo-forwarding-fragments-01, February 2014.

[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.

[RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", RFC 7102, January 2014.

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