The IPv6 Flow Label within a RPL domain
draft-thubert-6man-flow-label-for-rpl-01

Abstract

   This document present how the Flow Label can be used inside a RPL
   domain 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 saves 44 bits in each frame, and an eventual IP-
   in-IP encapsulation within the RPL domain that is required 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 November 14, 2014.

Copyright Notice

Table of Contents

1. Introduction

The emergence of radio technology enabled a large variety of new
types of devices to be interconnected, at a very low marginal cost
compared to wire, at any range from Near Field to interplanetary
distances, and in circumstances where wiring would be less than
practical, for instance rotating devices.

In particular, IEEE802.14.5 [IEEE802154] that is chartered to specify
PHY and MAC layers for radio Lowpower Lossy Networks (LLNs), defined
the TimeSlotted Channel Hopping [I-D.ietf-6tisch-tsch] (TSCH) mode of
operation as part of the IEEE802.15.4e MAC specification in order to
address Time Sensitive applications.

The 6TISCH architecture [I-D.ietf-6tisch-architecture] specifies the
operation IPv6 over TSCH wireless networks attached and synchronized
by backbone routers. In that model, route Computation may be
achieved in a centralized fashion by a Path Computation Element
(PCE), in a distributed fashion using the Routing Protocol for Low
Power and Lossy Networks [RFC6550] (RPL), or in a mixed mode. The
Backbone Routers may typically serve as roots for the RPL domain.

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. ISA100.11a
[ISA100.11a] is an example of such industrial WSN standard, using
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 SDO will scrutinize any bit in the frame and
reject any waste.

The challenge to obtain the adoption of IPv6 in the original standard
was really to save any possible bit 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.

1.1. On Wasted Energy

The design of Lowpower Lossy Networks is generally focussed on saving
energy, which is the most constrained resource of all. The other
constraints, such as the memory capacity and the duty cycling of the
LLN devices, derive from that primary concern. 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 Routing Protocol for Low Power and Lossy Networks (RPL) [RFC6550]
specification defines a generic Distance Vector protocol that is
indeed designed for very low energy consumption and 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 and set the related RPL information in
the packets.

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. In particular, the Rank, which is
the scalar metric computed by an specialized Objective Function such
as [RFC6552], is modified at each hop and allows to validate that the
packet progresses in the expected direction each upwards or downwards
in along the DODAG.

With [RFC6553] the RPL option is encoded as 6 Octets; it must be
placed in a Hop-by-Hop header that represents 2 additional octets for

a total of 8. In order to limit its range to the inside the RPL
domain, the Hop-by-Hop header 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 to the source, this operation
involves an extra IP-in-IP encapsulation inside the RPL domain for
all the packets which path is not contained within the RPL domain.

The 8-octets overhead 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-6lo-forwarding-fragments].

```
            ------+---------                          ^
                  |          Internet                 |
                  |                                   | Native IPv6
            +-----+                                   |
            |     | Border Router (RPL Root)    ^     |    ^
            |     |                             |     |    |
            +-----+                             |     |    | IPv6 +
                  |                             |     |    | HbH
           o    o    o     o                    |     |    | headers
         o o   o  o    o  o  o o    o            |     |    |
        o  o o  o o      o   o    o  o  o        |     |    |
         o   o    o  o      o  o    o  o  o       |     |    |
        o  o   o  o   o          o    o o        v     v    v
         o        o  o           o          o o
         o          o              o      o
```

                     LLN
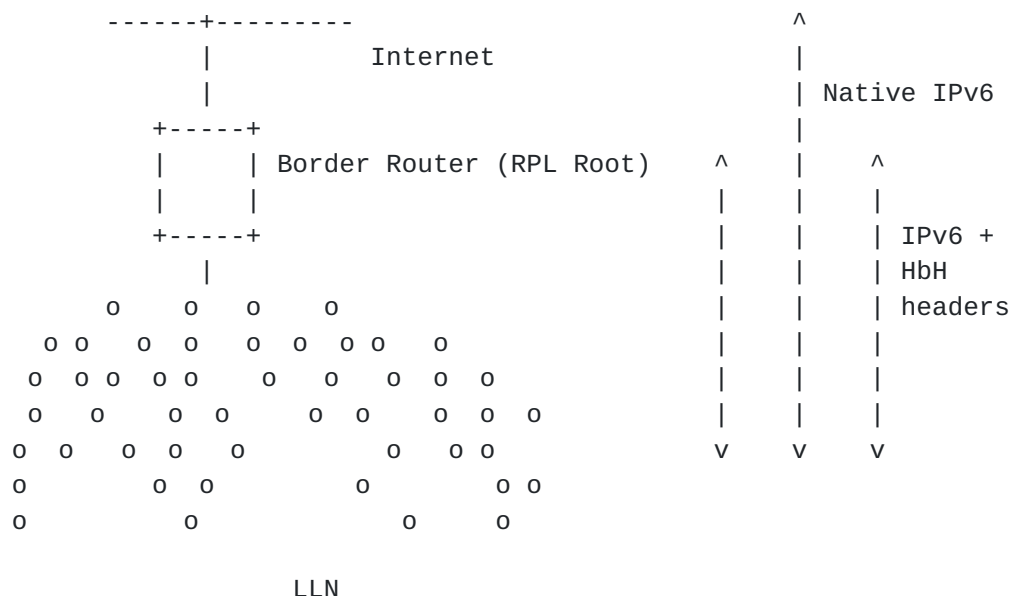
            Figure 1: IP-in-IP Encapsulation within the LLN

Considering that, in the classical IEEE802.14.5 PHY that is used by
ISA100.11a, roughly 80 octets are available per frame after security
is applied, and any additional transmitted bit weights in the energy
consumption and drains the batteries.

Regrettably, [RFC6282] does 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 IP-in-IP encapsulation, the
use of the flow label instead of the RPL option would be a valuable
saving.

1.2. 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
   RPL domain, is centralized to the root of the RPL DODAG structure, 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 from the Internet over the
   constrained network to a destination node that has no use of it.

1.3. 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.

   On the other hand, the operation of resetting or reusing the IPv6
   Flow Label at the root of a RPL domain is a deviation from the IPv6
   Flow Label Specification [RFC6437], in that it is neither the source
   nor the first hop router that sets the final Flow Label for use
   outside the RPL domain.

   Additionally, using the Flow Label to transport the information that
   is classically present in the RPL option implies that the Flow Label
   is modified at each hop inside the RPL domain, which again is a
   limited deviation from [RFC6437], which explicitly requires that the
   flow label cannot be modified once set.

   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 root 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.

   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 overhead for transporting the RPL
   information can be down from 64 to 20 bits, alleviating at the same
   time the need for IP-in-IP encapsulation. This optimization cannot
   be ignored, and can make the difference for the adoption of RPL and
   6TiSCH by external standard bodies.

   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' [RFC7102] and [RFC6550].

3. Applicability

This specification applies to a RPL [RFC6282] domain that forms a
stub LLN and is connected to the Internet by and only by its RPL
root(s), which act(s) as Border Router(s) for the LLN. With RPL, a
root is the bottleneck for all the traffic between the Internet and
the Destination-Oriented Directed Acyclic Graph (DODAG) that it
serves.

In that context, the specification entitles a RPL root to rewrite the
IPv6 [RFC2460] Flow Label of all packets entering or leaving the RPL
domain in both directions, from and towards the Internet, regardless
of its original setting. This may seem contradictory with 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, and
that security is a case justifying such a violation. This
specification suggests that energy-saving is another compelling
reason for a violation to the aforementioned rule.

For the compelling reason of saving energy, this specification allows
that regardless of its original setting, a root of a RPL domain MAY
reset the Flow Label of IPv6 packets entering the RPL domain to zero
for an optimal Header Compression by 6LoWPAN [RFC6282]. The
specification also allows that the root and LLN routers MAY reuse the
Flow Label inside the LLN for LLN purposes, such as to carry the RPL
Information as detailed hereafter.

This specification also allows that regardless of its original
setting, a a root of a RPL domain MAY set low Label of IPv6 packets
that exits the RPL domain MAY be set by the RPL, in a manner that
SHOULD conform the prescriptions in [RFC6437], and that a source in
the RPL domain MAY NOT expect that its setting of the Flow Label be
preserved end-to-end. From there, the capability by RPL routers
inside the LLN to alter a non-zero Flow Label between the source and
the root is another minor deviation to [RFC6437] that is also
acceptable since it is transparent to the core of the Internet.

4. **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:

      DAGRank(rank) = floor(rank/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.

5. **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 the
   property of belonging to a statistically uniform distribution that is
   desirable in the rest of the Internet.

   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.

5.1. 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.

5.2. 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.

6. 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".

7. Security Considerations

Because the flow label is not protected by IPSec, it is expected that
Layer-2 security is deployed in the LLN where is 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.

If the link layer is secured adequately, using the Flow Label as
opposed to the RPL option does not create an opening for a new threat
compared to [RFC6553].

8. IANA Considerations

No IANA action is required for this specification.

9. Acknowledgements

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

[10](#). References

[10.1](#). 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, "ISA100, Wireless Systems for Automation", May 2008,
          < 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. and R. Hinden, "Internet Protocol, Version 6
          (IPv6) Specification", RFC 2460, December 1998.

[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.

[10.2](#). 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", draft-ietf-6tisch-architecture-01 (work in
          progress), 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", draft-ietf-6tisch-tsch-00 (work in
          progress), November 2013.

[I-D.thubert-6lo-forwarding-fragments]
          Thubert, P. and J. Hui, "LLN Fragment Forwarding and
          Recovery", draft-thubert-6lo-forwarding-fragments-01 (work
          in progress), 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