

ROLL
Internet-Draft
Updates: [6550](#), [6553](#), [8138](#) (if approved)
Intended status: Standards Track
Expires: November 25, 2019

P. Thubert, Ed.
Cisco Systems
R. Jadhav
Huawei Tech
M. Gillmore
Itron
J. Pylakutty
Cisco
May 24, 2019

**Root initiated routing state in RPL
draft-ietf-roll-dao-projection-06**

Abstract

This document extends [RFC 6550](#), [RFC 6553](#) and [RFC 8138](#) and enable to install a limited amount of centrally-computed routes in a RPL graph, enabling loose source routing down a non-storing mode DODAG, or transversal routes inside the DODAG. In constrast with classical routes in RPL that are injected by the end devices, this draft enables the root of the DODAG to projects the routes that are needed on the nodes where they should be installed.

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 <https://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 25, 2019.

Copyright Notice

Copyright (c) 2019 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

(<https://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. BCP 14](#) [4](#)
 - [2.2. New Terms](#) [4](#)
 - [2.3. References](#) [4](#)
- [3. Extending RFC 6550](#) [4](#)
 - [3.1. RPL Instances](#) [5](#)
 - [3.2. New RPL Control Message Options](#) [5](#)
 - [3.3. RPI for Projected Routes](#) [7](#)
 - [3.4. Projected DAO](#) [7](#)
 - [3.4.1. Non-Storing Mode P-Route](#) [9](#)
 - [3.4.2. Storing-Mode P-Route](#) [10](#)
- [4. Extending RFC 8138](#) [13](#)
 - [4.1. Elective RPI 6LoRH](#) [13](#)
- [5. Extending RFC 6553](#) [13](#)
 - [5.1. Uncompressed RPL Option](#) [13](#)
- [6. Security Considerations](#) [14](#)
- [7. IANA Considerations](#) [14](#)
 - [7.1. New Elective 6LoWPAN Routing Header Type](#) [14](#)
 - [7.2. New RPL Control Codes](#) [15](#)
 - [7.3. Error in Projected Route ICMPv6 Code](#) [15](#)
- [8. Acknowledgments](#) [16](#)
- [9. References](#) [16](#)
 - [9.1. Normative References](#) [16](#)
 - [9.2. Informative References](#) [17](#)
- [Appendix A. Applications](#) [18](#)
 - [A.1. Loose Source Routing in Non-storing Mode](#) [18](#)
 - [A.2. Transversal Routes in storing and non-storing modes](#) [19](#)
- [Appendix B. Examples](#) [21](#)
 - [B.1. Using storing mode P-DAO in non-storing mode MOP](#) [21](#)
 - [B.2. Projecting a storing-mode transversal route](#) [22](#)
- [Authors' Addresses](#) [24](#)

1. Introduction

The "Routing Protocol for Low Power and Lossy Networks" [[RFC6550](#)] (LLN)(RPL) is a generic Distance Vector protocol that is well suited low energy Internet of Things (IoT) networks. RPL forms Destination

Oriented Directed Acyclic Graphs (DODAGs) in which the 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.

The 6TiSCH architecture [[I-D.ietf-6tisch-architecture](#)] leverages RPL for its routing operation and considers the Deterministic Networking Architecture [[I-D.ietf-detnet-architecture](#)] as one possible model whereby the device resources and capabilities are exposed to an external controller which installs routing states into the network based on some objective functions that reside in that external entity.

Based on heuristics of usage, path length, and knowledge of device capacity and available resources such as battery levels and reservable buffers, a Path Computation Element ([\[PCE\]](#)) with a global visibility on the system could install additional P2P routes that are more optimized for the current needs as expressed by the objective function.

This draft enables a RPL root to install and maintain projected routes (P-Routes) within its DODAG, along a selected set of nodes that may or may not include self, for a chosen duration. This potentially enables routes that are more optimized than those obtained with the distributed operation of RPL, either in terms of the size of a source-route header or in terms of path length, which impacts both the latency and the packet delivery ratio. P-routes may be installed in either Storing and Non-Storing Modes Instances of the classical RPL operation, resulting in potentially hybrid situations where the mode of some P-routes is different from that of the other routes in the RPL Instance.

P-Routes must be used with the parsimony to limit the amount of state that is installed in each device to fit within its resources, and to limit the amount of rerouted traffic to fit within the capabilities of the transmission links. The algorithm used to compute the paths and the protocol used to learn the topology of the network and the resources that are available in devices and in the network are out of scope for this document. Possibly with the assistance of a Path Computation Element ([\[PCE\]](#)) that could have a better visibility on the larger system, the root computes which segment could be optimized and uses this draft to install the corresponding P-Routes.

2. Terminology

2.1. BCP 14

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)][[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2.2. New Terms

P-Route: A route that is installed remotely by a RPL root.

2.3. References

In this document, readers will encounter terms and concepts that are discussed in the following documents:

- o "Routing Protocol for Low Power and Lossy Networks" [[RFC6550](#)], and
- o "Terminology in Low power And Lossy Networks" [[RFC7102](#)].

3. Extending RFC 6550

[Section 6.7](#) of RPL [[RFC6550](#)] specifies Control Message Options (CMO) to be placed in RPL messages such as the Destination Advertisement Object (DAO) message. The RPL Target Option and the Transit Information Option (TIO) are such options. In Non-Storing Mode, the TIO option is used in the DAO message to indicate the immediate parent of a given path. The TIO applies to the Target options that immediately precede it. Options may be factorized; multiple TIOs may be present to indicate multiple routes to the one or more contiguous addresses indicated in the Target Options that immediately precede the TIOs in the RPL message.

This specification introduces two new Control Message Options referred to as Route Projection Options (RPO). One RPO is the Information option (VIO) and the other is the Source-Routed VIO (SRVIO). The VIO installs a route on each hop along a P-Route (in a fashion analogous to RPL Storing Mode) whereas the SRVIO installs a source-routing state at the ingress node, which uses it to insert a routing header in a fashion similar to Non-Storing Mode.

Like the TIO, the RPOs MUST be preceded by one or more RPL Target Options to which they apply, and they can be factorized: multiple contiguous RPOs indicate alternate paths to the target(s).

3.1. RPL Instances

It must be noted that RPL has a concept of instance but does not have a concept of an administrative distance, which exists in certain proprietary implementations to sort out conflicts between multiple sources of routing information. This draft conforms the instance model as follows:

- o If the PCE needs to influence a particular instance to add better routes in conformance with the routing objectives in that instance, it may do so. When the PCE modifies an existing instance then the added routes must not create a loop in that instance. This is achieved by always preferring a route obtained from the PCE over a route that is learned via RPL.
- o If the PCE installs a more specific (say, Traffic Engineered) route between a particular pair of nodes then it SHOULD use a Local Instance from the ingress node of that path. A packet associated with that instance will be routed along that path and MUST NOT be placed over a Global Instance again. A packet that is placed on a Global Instance may be injected in the Local Instance based on node policy and the Local Instance parameters.

In all cases, the path is indicated by a new Via Information option, and the flow is similar to the flow used to obtain loose source routing.

3.2. New RPL Control Message Options

The format of RPOs is as follows:

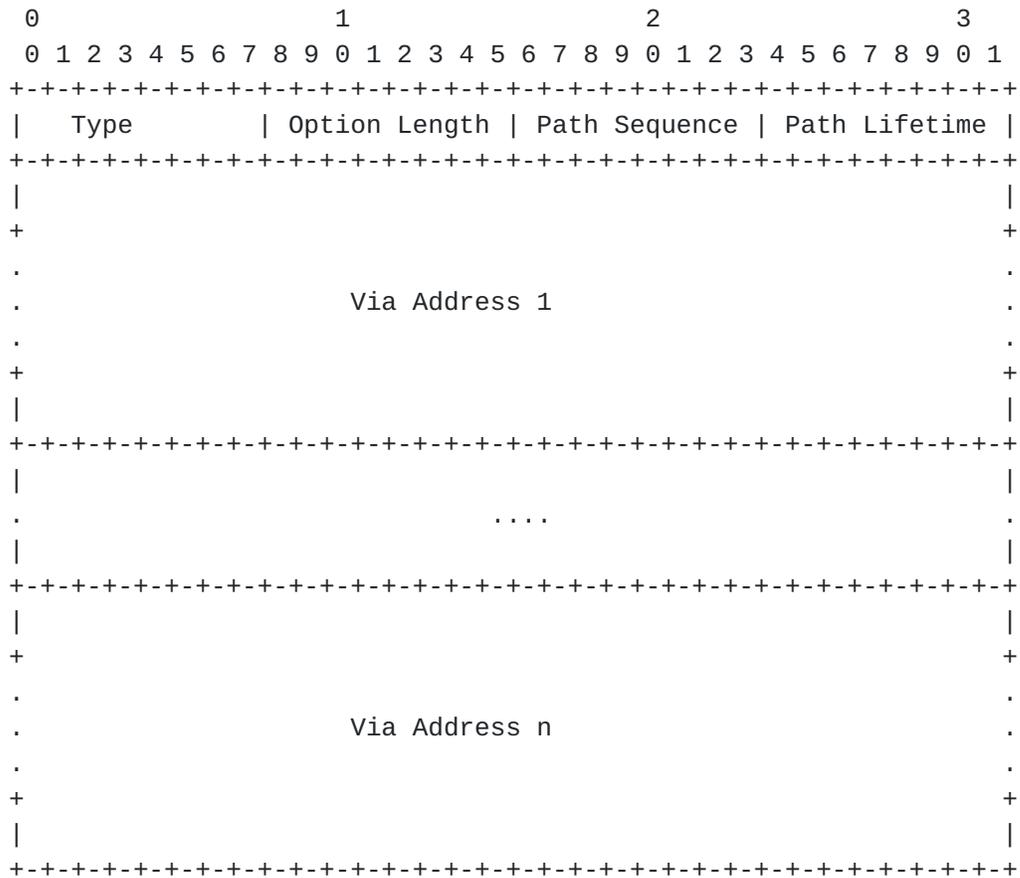


Figure 1: Via Information option format

Option Type: 0x0A for VIO, 0x0B for SRVIO (to be confirmed by IANA)

Option Length: In bytes; variable, depending on the number of Via Addresses.

Path Sequence: 8-bit unsigned integer. When a RPL Target option is issued by the root of the DODAG (i.e. in a DAO message), that root sets the Path Sequence and increments the Path Sequence each time it issues a RPL Target option with updated information. The indicated sequence deprecates any state for a given Target that was learned from a previous sequence and adds to any state that was learned for that sequence.

Path Lifetime: 8-bit unsigned integer. The length of time in Lifetime Units (obtained from the Configuration option) that the prefix is valid for route determination. The period starts when a new Path Sequence is seen. A value of 255 (0xFF) represents infinity. A value of zero (0x00) indicates a loss of reachability. A DAO message that contains a Via Information

option with a Path Lifetime of zero for a Target is referred as a No-Path (for that Target) in this document.

Via Address: 16 bytes. IPv6 Address of the next hop towards the destination(s) indicated in the target option that immediately precede the RPO. Via Addresses are indicated in the order of the data path from the ingress to the egress nodes.

An RPO MUST contain at least one Via Address, and a Via Address MUST NOT be present more than once, otherwise the RPO MUST be ignored.

3.3. RPI for Projected Routes

RPL [\[RFC6550\]](#), [Section 11.2](#), specifies the RPL Packet Information (RPI) as a set of fields that are placed by RPL routers in IP packets to identify the RPL Instance, detect anomalies and trigger corrective actions.

In particular, the SenderRank, which is the scalar metric computed by a specialized Objective Function such as described in [\[RFC6552\]](#), indicates the Rank of the sender and is modified at each hop. The SenderRank field is used to validate that the packet progresses in the expected direction, either upwards or downwards, along the DODAG.

RPL defines the "RPL Option for Carrying RPL Information in Data-Plane Datagrams" [\[RFC6553\]](#) to transport the RPI, which is carried in an IPv6 Hop-by-Hop Options Header [\[RFC8200\]](#), typically consuming eight bytes per packet.

This specification updates [\[RFC6550\]](#) as follows. When using projected routes, the Rank is useless and SHOULD be set to 0 in the non-compressed form, and can be elided in the compressed form (see [Section 4.1](#)). In a same fashion, the O, R, and F flags that are defined in [Section 11.2 of \[RFC6550\]](#) are not used for packets that follow a projected route and they MUST be reset. A new flag is added, the P flag that indicates that the packet is injected along a projected route.

3.4. Projected DAO

This draft adds a capability to RPL whereby the root of a DODAG projects a route by sending an extended DAO message called a Projected-DAO (P-DAO) to an arbitrary router in the DODAG, indicating one or more sequence(s) of routers inside the DODAG via which the target(s) indicated in the Target Information Option(s) (TIO) can be reached.

A P-DAO is sent from a global address of the root to a global address of the recipient, and MUST be confirmed by a DAO-ACK, which is sent back to a global address of the root.

A P-DAO message MUST contain at least one TIO and at least one RPO following it. There can be at most one such sequence of TIOs and then RPOs.

Like a classical DAO message, a P-DAO is processed only if it is "new" per [section 9.2.2](#). "Generation of DAO Messages" of the RPL specification [[RFC6550](#)]; this is determined using the Path Sequence information from the RPO as opposed to a TIO. Also, a Path Lifetime of 0 in an RPO indicates that a route is to be removed.

There are two kinds of operation for the P-Routes, the Storing Mode and the Non-Storing Mode.

- o The Non-Storing Mode is discussed in [Section 3.4.1](#). It uses an SRVIO that carries a list of Via Addresses to be used as a source-routed path to the target. The recipient of the P-DAO is the ingress router of the source-routed path. Upon a Non-Storing Mode P-DAO, the ingress router installs a source-routed state to the target and replies to the root directly with a DAO-ACK message.
- o The Storing Mode is discussed in [Section 3.4.2](#). It uses a VIO with one Via Address per consecutive hop, from the ingress to the egress of the path, including the list of all intermediate routers in the data path order. The Via Addresses indicate the routers in which the routing state to the target have to be installed via the next Via Address in the VIO. In normal operations, the P-DAO is propagated along the chain of Via Routers from the egress router of the path till the ingress one, which confirms the installation to the root with a DAO-ACK message. Note that the root may be the ingress and it may be the egress of the path, that it can also be neither but it cannot be both.

In case of a forwarding error along a P-Route, an ICMP error is sent to the root with a new Code "Error in Projected Route" (See [Section 7.3](#)). The root can then modify or remove the P-Route. The "Error in Projected Route" message has the same format as the "Destination Unreachable Message", as specified in [RFC 4443](#) [[RFC4443](#)]. The portion of the invoking packet that is sent back in the ICMP message SHOULD record at least up to the routing header if one is present, and the routing header SHOULD be consumed by this node so that the destination in the IPv6 header is the next hop that this node could not reach. if a 6LoWPAN Routing Header (6LoRH) [[RFC8138](#)] is used to carry the IPv6 routing information in the outer header then that whole 6LoRH information SHOULD be present in the

case of a loose source-routed path, there MUST be either a neighbor that is adjacent to the loose next hop, on which case the packet is forwarded to that neighbor, or a source-routed path to the loose next hop; in the latter case, another encapsulation takes place and the process possibly recurses; otherwise the packet is dropped.

In order to add a source-routing header, the router encapsulates the packet with an IP-in-IP header and a non-storing mode source routing header (SRH) [RFC6554]. In the uncompressed form the source of the packet would be self, the destination would be the first Via Address in the SRVIO, and the SRH would contain the list of the remaining Via Addresses and then the target.

In practice, the router will normally use the "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Paging Dispatch" [RFC8025] to compress the RPL artifacts as indicated in the "6LoWPAN Routing Header" [RFC8138] specification. In that case, the router indicates self as encapsulator in an IP-in-IP 6LoRH Header, and places the list of Via Addresses in the order of the VIO and then the target in the SRH 6LoRH Header.

```

+-+ ... -+-+ ... +-+ ... -+-+ ...      +-+ ... -+-+ ...
|11110001|SRH-6LoRH| ERPI- | IP-in-IP  Encap | NH=1      |11110CPP|
|Page 1  |Type1 S=2| 6LoRH | 6LoRH  sulator |LOWPAN_IPHC| UDP    |
+-+ ... -+-+ ... +-+ ... -+-+ ...      +-+ ... -+-+ ...
      <-RFC8138-><-This-><----RFC 8138----><-----RFC 6282----->
                    RFC          5 to 19 bytes      No RPL artifact
    
```

Figure 3: Example Compressed Packet with SRH.

In case of a forwarding error along a Source Route path, the node that fails to forward SHOULD send an ICMP error with a code "Error in Source Routing Header" back to the source of the packet, as described in section 11.2.2.3. of [RFC6550]. Upon this message, the encapsulating node SHOULD stop using the source route path for a period of time and it SHOULD send an ICMP message with a Code "Error in Projected Route" to the root. Failure to follow these steps may result in packet loss and wasted resources along the source route path that is broken.

3.4.2. Storing-Mode P-Route

As illustrated in Figure 4, the Storing Mode projected iq used by the root to install a routing state towards a target in the routers along a segment between an ingress and an egress router; this enables the routers to forward along that segment any packet for which the next loose hop is the said target, for instance a loose source routed packet for which the next loose hop is the target, or a packet for

which the router has a routing state to the final destination via the target.

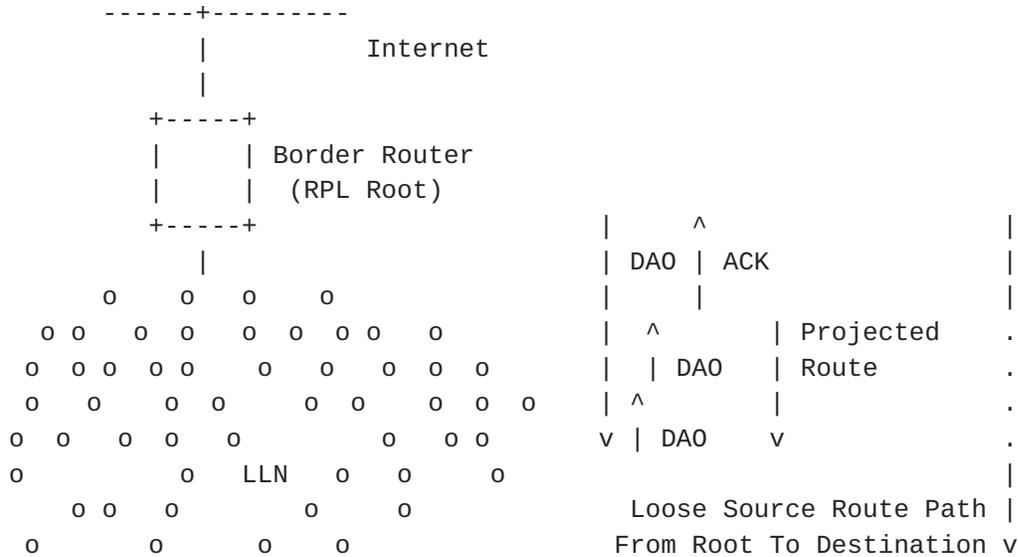


Figure 4: Projecting a route

In order to install the relevant routing state along the segment between an ingress and an egress routers, the root sends a unicast P-DAO message to the egress router of the routing segment that must be installed. The P-DAO message contains the ordered list of hops along the segment as a direct sequence of Via Information options that are preceded by one or more RPL Target options to which they relate. Each Via Information option contains a Path Lifetime for which the state is to be maintained.

The root sends the P-DAO directly to the egress node of the segment. In that P-DAO, the destination IP address matches the Via Address in the last VIO. This is how the egress recognizes its role. In a similar fashion, the ingress node recognizes its role as it matches Via Address in the first VIO.

The egress node of the segment is the only node in the path that does not install a route in response to the P-DAO; it is expected to be already able to route to the target(s) on its own. It may either be the target, or may have some existing information to reach the target(s), such as a connected route or an already installed P-Route. If one of the targets cannot be located, the node MUST answer to the root with a negative DAO-ACK listing the target(s) that could not be located (suggested status 10 to be confirmed by IANA).

If the egress node can reach all the targets, then it forwards the P-DAO with unchanged content to its loose predecessor in the segment as indicated in the list of Via Information options, and recursively the message is propagated unchanged along the sequence of routers indicated in the P-DAO, but in the reverse order, from egress to ingress.

The address of the predecessor to be used as destination of the propagated DAO message is found in the Via Information option the precedes the one that contain the address of the propagating node, which is used as source of the packet.

Upon receiving a propagated DAO, an intermediate router as well as the ingress router install a route towards the DAO target(s) via its successor in the P-DAO; the router locates the VIO that contains its address, and uses as next hop the address found in the Via Address field in the following VIO. The router MAY install additional routes towards the addresses that are located in VIOs that are after the next one, if any, but in case of a conflict or a lack of resource, a route to a target installed by the root has precedence.

The process recurses till the P-DAO is propagated to ingress router of the segment, which answers with a DAO-ACK to the root.

Also, the path indicated in a P-DAO may be loose, in which case the reachability to the next hop has to be asserted. Each router along the path indicated in a P-DAO is expected to be able to reach its successor, either with a connected route (direct neighbor), or by routing, for instance following a route installed previously by a DAO or a P-DAO message. If that route is not connected then a recursive lookup may take place at packet forwarding time to find the next hop to reach the target(s). If it does not and cannot reach the next router in the P-DAO, the router MUST answer to the root with a negative DAO-ACK indicating the successor that is unreachable (suggested status 11 to be confirmed by IANA).

A Path Lifetime of 0 in a Via Information option is used to clean up the state. The P-DAO is forwarded as described above, but the DAO is interpreted as a No-Path DAO and results in cleaning up existing state as opposed to refreshing an existing one or installing a new one.

In case of a forwarding error along a Storing Mode P-Route, the node that fails to forward SHOULD send an ICMP error with a code "Error in Projected Route" to the root. Failure to do so may result in packet loss and wasted resources along the P-Route that is broken.

4. Extending RFC 8138

4.1. Elective RPI 6LoRH

[RFC8138] defines a Critical 6LoRH to compress the RPL RPI found in normal packets inside a RPL domain, the RPI-6LoRH.

this specification introduces the ERPI-6LoRH header that MUST be used to compress the RPI in packets that follow a projected route. As discussed in Section 3.3, the Rank and the O, R, and F flags are always set to 0 and can be elided. The new P flag is always set and can also be elided. It results that in general only the RPL InstanceID is necessary in the compressed form.

This specification adds an optimization whereby the local RPLInstanceID 0 for the source of the packet (the encapsulator when using IP in IP) can be elided. This is the case where the RPLInstanceID is encoded as binary b10000000, decimal 128, in the non-compressed form.

The ERPI-6LoRH header is Elective since it does not contain information that is critical to the routing and it can be ignored when not understood. The resulting format is illustrated in Figure 5 below:

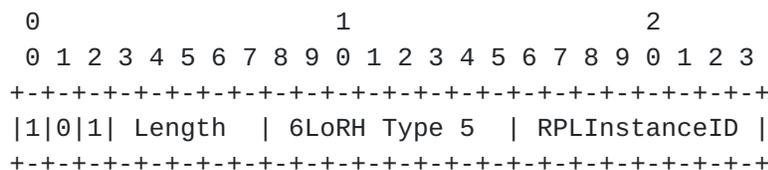


Figure 5: A ERPI-6LoRH carrying a RPLInstanceID

The ERPI-6LoRH header is identifies by a 6LoRH Type of 5 (to be confirmed by IANA), which is the same value as the RPI-6LoRH but in the Elective namespace.If the RPLInstanceID is a local RPLInstanceID 0 for the source of the packet then it MUST be elided and the length MUST be set to 0. Else the length MUST be set to 1 to indicate that the ERPI-6LoRH carries a RPLInstanceID.

5. Extending RFC 6553

5.1. Uncompressed RPL Option

[RFC6553] defines a format for the RPI that is suitable for transporting in the IPv6 Hop-by-Hop Header [RFC8200]. This

specification introduces a new flag in the RPI that must be encoded in any format includeing uncompressed.

The updated format for the RPL Option is presented in Figure 6.

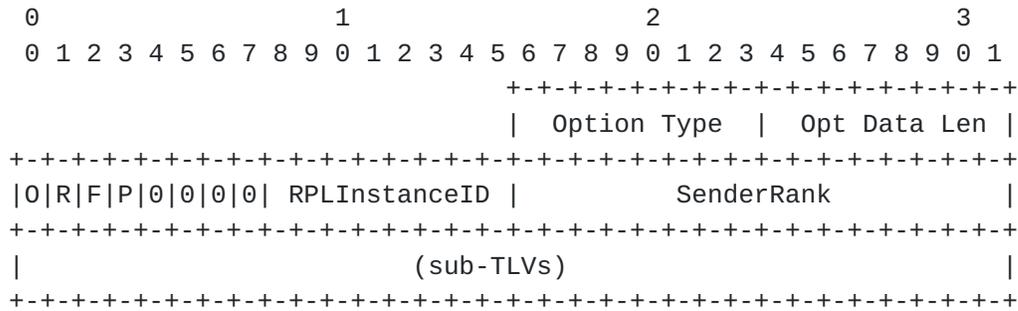


Figure 6: RPL Option

New fields:

P: 1-bit flag; indicates that the packet is routed along a projected route.

6. Security Considerations

This draft uses messages that are already present in RPL [RFC6550] with optional secured versions. The same secured versions may be used with this draft, and whatever security is deployed for a given network also applies to the flows in this draft.

TODO: should probably consider how P-DAO messages could be abused by a) rogue nodes b) via replay of messages c) if use of P-DAO messages could in fact deal with any threats?

7. IANA Considerations

7.1. New Elective 6LoWPAN Routing Header Type

This specification assigns a new value (to be confirmed by IANA) in the Elective 6LoWPAN Routing Header Type Registry created for RFC 8138 as below:

Value	Meaning	Defining Spec
5 (suggested)	ERPI-6LoRH	This document

Table 1: New Elective 6LoWPAN Routing Header Type

7.2. New RPL Control Codes

This document extends the IANA registry created by [RFC 6550](#) for RPL Control Codes as follows:

Code	Description	Reference
0x0A	Via	This document
0x0B	Source-Routed Via	This document

RPL Control Codes

This document is updating the registry created by [RFC 6550](#) for the RPL 3-bit Mode of Operation (MOP) as follows:

MOP value	Description	Reference
5	Non-Storing mode of operation with P-Routes	This document
6	Storing mode of operation with P-Routes	This document

DIO Mode of operation

7.3. Error in Projected Route ICMPv6 Code

In some cases RPL will return an ICMPv6 error message when a message cannot be forwarded along a P-Route. This ICMPv6 error message is "Error in Projected Route".

IANA has defined an ICMPv6 "Code" Fields Registry for ICMPv6 Message Types. ICMPv6 Message Type 1 describes "Destination Unreachable" codes. This specification requires that a new code is allocated from the ICMPv6 Code Fields Registry for ICMPv6 Message Type 1, for "Error

in Projected Route", with a suggested code value of 8, to be confirmed by IANA.

8. Acknowledgments

The authors wish to acknowledge JP Vasseur and Patrick Wetterwald for their contributions to the ideas developed here.

9. References

9.1. 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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, [RFC 4443](#), DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., 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](#), DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC6552] Thubert, P., Ed., "Objective Function Zero for the Routing Protocol for Low-Power and Lossy Networks (RPL)", [RFC 6552](#), DOI 10.17487/RFC6552, March 2012, <<https://www.rfc-editor.org/info/rfc6552>>.
- [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](#), DOI 10.17487/RFC6553, March 2012, <<https://www.rfc-editor.org/info/rfc6553>>.
- [RFC6554] Hui, J., Vasseur, JP., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL)", [RFC 6554](#), DOI 10.17487/RFC6554, March 2012, <<https://www.rfc-editor.org/info/rfc6554>>.

- [RFC8025] Thubert, P., Ed. and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Paging Dispatch", [RFC 8025](#), DOI 10.17487/RFC8025, November 2016, <<https://www.rfc-editor.org/info/rfc8025>>.
- [RFC8138] Thubert, P., Ed., Bormann, C., Toutain, L., and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing Header", [RFC 8138](#), DOI 10.17487/RFC8138, April 2017, <<https://www.rfc-editor.org/info/rfc8138>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, [RFC 8200](#), DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

9.2. Informative References

- [I-D.ietf-6tisch-architecture]
Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", [draft-ietf-6tisch-architecture-20](#) (work in progress), March 2019.
- [I-D.ietf-detnet-architecture]
Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", [draft-ietf-detnet-architecture-13](#) (work in progress), May 2019.
- [PCE] IETF, "Path Computation Element", <<https://datatracker.ietf.org/doc/charter-ietf-pce/>>.
- [RFC6997] Goyal, M., Ed., Baccelli, E., Philipp, M., Brandt, A., and J. Martocci, "Reactive Discovery of Point-to-Point Routes in Low-Power and Lossy Networks", [RFC 6997](#), DOI 10.17487/RFC6997, August 2013, <<https://www.rfc-editor.org/info/rfc6997>>.
- [RFC7102] Vasseur, JP., "Terms Used in Routing for Low-Power and Lossy Networks", [RFC 7102](#), DOI 10.17487/RFC7102, January 2014, <<https://www.rfc-editor.org/info/rfc7102>>.

As a network gets deep, the size of the source routing header that the root must add to all the downward packets becomes an issue for nodes that are many hops away. In some use cases, a RPL network forms long lines and a limited amount of well-targeted routing state would allow to make the source routing operation loose as opposed to strict, and save packet size. Limiting the packet size is directly beneficial to the energy budget, but, mostly, it reduces the chances of frame loss and/or packet fragmentation, which is highly detrimental to the LLN operation. Because the capability to store a routing state in every node is limited, the decision of which route is installed where can only be optimized with a global knowledge of the system, a knowledge that the root or an associated PCE may possess by means that are outside of the scope of this specification.

This specification enables to store source-routed or storing mode state in intermediate routers, which enables to limit the excursion of the source route headers in deep networks. Once a P-DAO exchange has taken place for a given target, if the root operates in non storing mode, then it may elide the sequence of routers that is installed in the network from its source route headers to destination that are reachable via that target, and the source route headers effectively become loose.

A.2. Transversal Routes in storing and non-storing modes

RPL is optimized for Point-to-Multipoint (P2MP) and Multipoint-to-Point (MP2P), whereby routes are always installed along the RPL DODAG respectively from and towards the DODAG Root. Transversal Peer to Peer (P2P) routes in a RPL network will generally suffer from some elongated (stretched) path versus the best possible path, since routing between 2 nodes always happens via a common parent, as illustrated in Figure 8:

- o in non-storing mode, all packets routed within the DODAG flow all the way up to the root of the DODAG. If the destination is in the same DODAG, the root must encapsulate the packet to place a Routing Header that has the strict source route information down the DODAG to the destination. This will be the case even if the destination is relatively close to the source and the root is relatively far off.
- o In storing mode, unless the destination is a child of the source, the packets will follow the default route up the DODAG as well. If the destination is in the same DODAG, they will eventually reach a common parent that has a route to the destination; at worse, the common parent may also be the root. From that common parent, the packet will follow a path down the DODAG that is

optimized for the Objective Function that was used to build the DODAG.

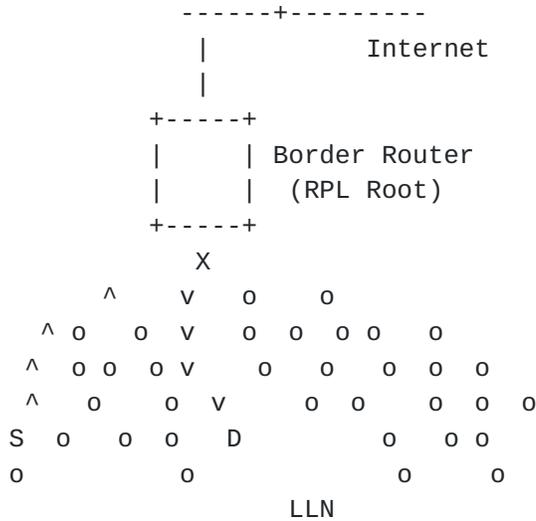


Figure 8: Routing Stretch between S and D via common parent X

It results that it is often beneficial to enable transversal P2P routes, either if the RPL route presents a stretch from shortest path, or if the new route is engineered with a different objective. For that reason, earlier work at the IETF introduced the "Reactive Discovery of Point-to-Point Routes in Low Power and Lossy Networks" [[RFC6997](#)], which specifies a distributed method for establishing optimized P2P routes. This draft proposes an alternate based on a centralized route computation.

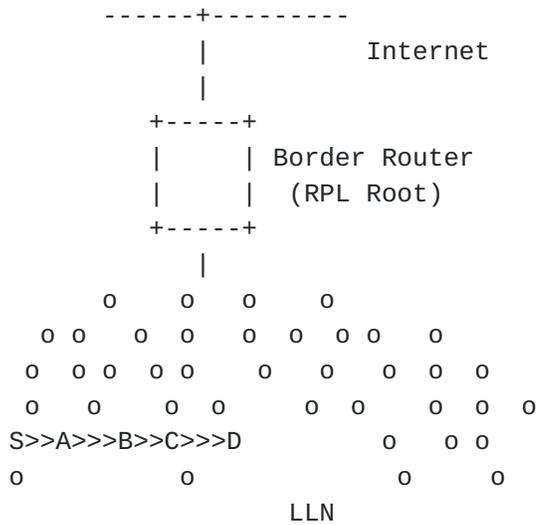


Figure 9: Projected Transversal Route

This specification enables to store source-routed or storing mode state in intermediate routers, which enables to limit the stretch of a P2P route and maintain the characteristics within a given SLA. An example of service using this mechanism could be a control loop that would be installed in a network that uses classical RPL for asynchronous data collection. In that case, the P2P path may be installed in a different RPL Instance, with a different objective function.

Appendix B. Examples

B.1. Using storing mode P-DAO in non-storing mode MOP

In non-storing mode, the DAG root maintains the knowledge of the whole DODAG topology, so when both the source and the destination of a packet are in the DODAG, the root can determine the common parent that would have been used in storing mode, and thus the list of nodes in the path between the common parent and the destination. For instance in the diagram shown in Figure 10, if the source is node 41 and the destination is node 52, then the common parent is node 22.

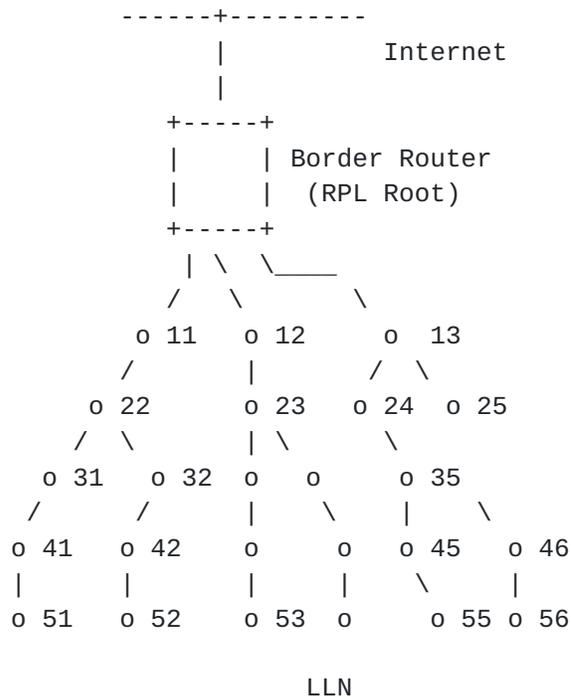


Figure 10: Example DODAG forming a logical tree topology

With this draft, the root can install a storing mode routing states along a segment that is either from itself to the destination, or from one or more common parents for a particular source/destination

pair towards that destination (in this particular example, this would be the segment made of nodes 22, 32, 42).

In the example below, say that there is a lot of traffic to nodes 55 and 56 and the root decides to reduce the size of routing headers to those destinations. The root can first send a DAO to node 45 indicating target 55 and a Via segment (35, 45), as well as another DAO to node 46 indicating target 56 and a Via segment (35, 46). This will save one entry in the routing header on both sides. The root may then send a DAO to node 35 indicating targets 55 and 56 a Via segment (13, 24, 35) to fully optimize that path.

Alternatively, the root may send a DAO to node 45 indicating target 55 and a Via segment (13, 24, 35, 45) and then a DAO to node 46 indicating target 56 and a Via segment (13, 24, 35, 46), indicating the same DAO Sequence.

B.2. Projecting a storing-mode transversal route

In this example, say that a PCE determines that a path must be installed between node S and node D via routers A, B and C, in order to serve the needs of a particular application.

The root sends a P-DAO with a target option indicating the destination D and a sequence Via Information option, one for S, which is the ingress router of the segment, one for A and then for B, which are an intermediate routers, and one for C, which is the egress router.

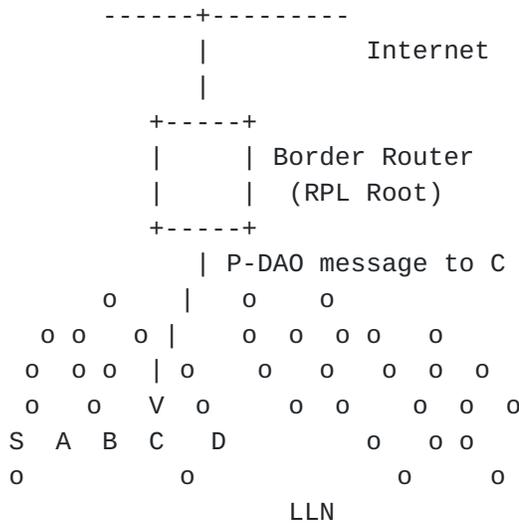


Figure 11: P-DAO from root

Upon reception of the P-DAO, C validates that it can reach D, e.g. using IPv6 Neighbor Discovery, and if so, propagates the P-DAO unchanged to B.

B checks that it can reach C and if so, installs a route towards D via C. Then it propagates the P-DAO to A.

The process recurses till the P-DAO reaches S, the ingress of the segment, which installs a route to D via A and sends a DAO-ACK to the root.

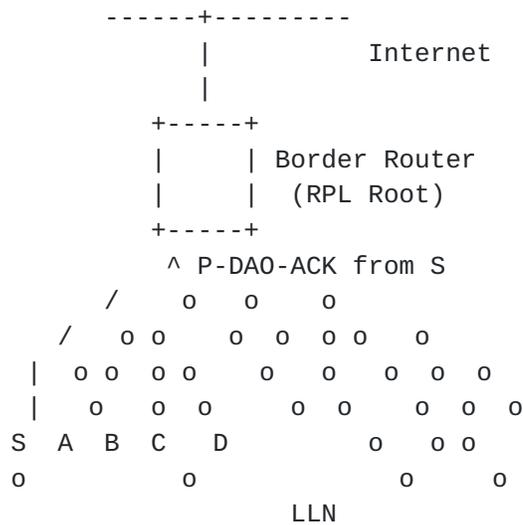


Figure 12: P-DAO-ACK to root

As a result, a transversal route is installed that does not need to follow the DODAG structure.

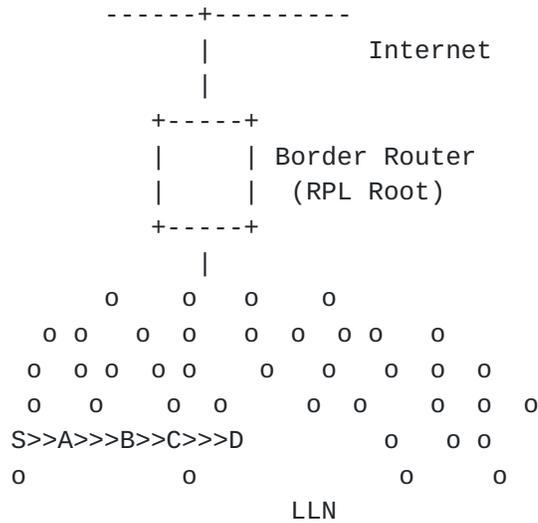


Figure 13: Projected Transversal Route

Authors' Addresses

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

Rahul Arvind Jadhav
 Huawei Tech
 Kundalahalli Village, Whitefield,
 Bangalore, Karnataka 560037
 India

Phone: +91-080-49160700
 Email: rahul.ietf@gmail.com

Matthew Gillmore
Itron, Inc
Building D
2111 N Molter Road
Liberty Lake 99019
United States

Phone: +1.800.635.5461
Email: matthew.gillmore@itron.com

James Pylakutty
Cisco Systems
Cessna Business Park
Kadubeesanahalli
Marathalli ORR
Bangalore, Karnataka 560087
INDIA

Phone: +91 80 4426 4140
Email: mundenma@cisco.com

