

Workgroup: ROLL

U [6550](#) (if approved)

p
d
a
t
e
s
:

Published: September 11, 2020

Intended Status: Standards Track

Expires: March 15, 2021

A P. Thubert, Ed. R.A. Jadhav M. Gillmore

uCisco Systems Huawei Tech Itron

t
h
o
r
s
:

Root initiated routing state in RPL

Abstract

This document enables a RPL Root to install and maintain Projected 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 or resilient than those obtained with the classical 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. Projected 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 Projected Routes is different from that of the other routes in the RPL Instance.

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 March 15, 2021.

Copyright Notice

Copyright (c) 2020 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. Terminology](#)
 - [2.1. Requirements Language](#)
 - [2.2. References](#)
 - [2.3. Glossary](#)
 - [2.4. Other Terms](#)
- [3. Updating RFC 6550](#)
- [4. Identifying a Path](#)
- [5. New RPL Control Messages and Options](#)
 - [5.1. New P-DAO Request Control Message](#)
 - [5.2. New PDR-ACK Control Message](#)
 - [5.3. Route Projection Options](#)
 - [5.4. Sibling Information Option](#)
- [6. Projected DAO](#)
 - [6.1. Non-Storing Mode Projected Route](#)
 - [6.2. Storing-Mode Projected Route](#)
- [7. Security Considerations](#)
- [8. IANA Considerations](#)
 - [8.1. New RPL Control Codes](#)
 - [8.2. New RPL Control Message Options](#)
 - [8.3. New SubRegistry for the Projected DAO Request \(PDR\) Flags](#)
 - [8.4. New SubRegistry for the PDR-ACK Flags](#)
 - [8.5. New Subregistry for the PDR-ACK Acceptance Status values](#)
 - [8.6. New Subregistry for the PDR-ACK Rejection Status values](#)
 - [8.7. New SubRegistry for the Route Projection Options \(RPO\) Flags](#)
 - [8.8. New SubRegistry for the Sibling Information Option \(SIO\) Flags](#)
 - [8.9. Error in Projected Route ICMPv6 Code](#)
- [9. Acknowledgments](#)
- [10. Normative References](#)
- [11. Informative References](#)
- [Appendix A. Applications](#)
 - [A.1. Loose Source Routing in Non-storing Mode](#)
 - [A.2. Transversal Routes in storing and non-storing modes](#)
- [Appendix B. Examples](#)
 - [B.1. Using storing mode P-DAO in non-storing mode MOP](#)
 - [B.2. Projecting a storing-mode transversal route](#)
- [Authors' Addresses](#)

1. Introduction

RPL, the ["Routing Protocol for Low Power and Lossy Networks"](#) [RPL] (LLNs), is a generic Distance Vector protocol that is well suited for application in a variety of 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"](#) [6TiSCH-ARCHI] uses RPL for its routing operations.

The 6TiSCH Architecture also leverages the ["Deterministic Networking Architecture"](#) [RFC8655] centralized 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. With DetNet and 6TiSCH, the component of the controller that is responsible of computing routes is called a Path Computation Element ([PCE]).

Based on heuristics of usage, path length, and knowledge of device capacity and available resources such as battery levels and reservable buffers, a PCE with a global visibility on the system can compute P2P routes that are more optimized for the current needs as expressed by the objective function.

This draft proposes protocol extensions to RPL that enable the Root to install a limited amount of centrally-computed routes in a RPL graph, on behalf of a PCE that may be collocated or separated from the Root. Those extensions enable loose source routing down in RPL Non-Storing Mode and transversal routes inside the DODAG regardless of the RPL Mode of Operation (MOP).

As opposed to the classical RPL operations where routes are injected by the Target nodes, the protocol extensions enable the Root of a DODAG to project the routes that are needed onto the nodes where they should be installed. This specification uses the term Projected Route to refer to those routes. A Projected Route may be a stand-alone end-to-end path to a Target or a segment in a more complex forwarding graph called a Track.

The concept of a Track was introduced in the 6TiSCH architecture, as a complex path with redundant forwarding solutions along the way. A node that is present on more than one segment in a Track may be able to use either of the Track segments to forward a packet towards the Target.

The ["Reliable and Available Wireless \(RAW\) Architecture/Framework"](#) [RAW-ARCHI] extends the the forward plane to leverage that redundancy with the Packet ARQ, Replication, Elimination, and Overhearing (PAREO) functions.

The RAW Architecture also discusses the dynamic selection of the best path for a packet within a Track at forwarding time. To that effect, it defines the Path Selection Engine (PSE), which is the counter-part of the PCE to perform rapid local adjustments of the

forwarding tables within the path diversity that the PCE has selected for the Track.

RAW differentiates the time scale at which the PCE computes the Track (slow, globally optimized, based on statistical metrics) and the time scale at which the PSE makes the forwarding decision for a packet or a small collection of packets (fast, but with a scope limited to the Track). This provides a dynamic balance between the reliability and availability requirements of the flows and the need to conserve energy and spectrum.

Projected Routes must be used with the parsimony to limit the amount of state that is installed in each device to fit within the device resources, and to maintain the amount of rerouted traffic within the capabilities of the transmission links. The methods used to learn the node capabilities and the resources that are available in the devices and in the network are out of scope for this document.

This specification expects that a base RPL Instance is operated in RPL Non-Storing Mode to sustain the exchanges with the Root. In that Mode, the Root has enough information to build a basic DODAG topology based on parents and children, but lacks the knowledge of siblings. This document adds the capability for nodes to advertise sibling information in order to improve the topological awareness of the Root.

This specification uses the RPL Root as a proxy to the PCE. The PCE may be collocated with the Root, or may reside in an external Controller. In that case, the PCE exchanges control messages with the Root over a Southbound API, that is out of scope for this specification. The algorithm to compute the paths and the protocol used by an external PCE to obtain the topology of the network from the Root are also out of scope.

2. Terminology

2.1. Requirements Language

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. References

In this document, readers will encounter terms and concepts that are discussed in the "[Routing Protocol for Low Power and Lossy Networks](#)" [[RPL](#)] and "[Terminology in Low power And Lossy Networks](#)" [[RFC7102](#)].

2.3. Glossary

This document often uses the following acronyms:

6BBR: 6LoWPAN Backbone Router

6LBR: 6LoWPAN Border Router

6LN:
6LowPAN Node

6LR: 6LowPAN Router

CMO: Control Message Option

DAO: Destination Advertisement Object

DODAG: Destination-Oriented Directed Acyclic Graph

LLN: Low-Power and Lossy Network

MOP: RPL Mode of Operation

NA: Neighbor Advertisement

NCE: Neighbor Cache Entry

ND: Neighbor Discovery

NDP: Neighbor Discovery Protocol

NS: Neighbor Solicitation

P-DAO: A Projected DAO is a DAO message sent by the RPL Root to install a Projected Route.

PDR P-DAO Request

RA: Router Advertisement

RAN: RPL-Aware Node

RS: Router Solicitation

RPL: IPv6 Routing Protocol for LLNs [[RPL](#)]

RPO: A Route Projection Option; it can be a VIO or an SRVIO.

RTO: RPL Target Option

SIO: RPL Sibling Information Option

SRVIO: A Source-Routed Via Information Option, used in Non-Storing Mode P-DAO messages.

TIO: RPL Transit Information Option

VIO: A Via Information Option, used in Storing Mode P-DAO messages.

2.4. Other Terms

Projected Route: A Projected Route is a serial path or path segment that is computed, installed and maintained remotely by a RPL Root.

Track: The term Track is used in this document to refer to a complex path, e.g., a DODAG, that incorporates redundant Projected

Routes towards a destination using diversity to increase the reliability.

3. Updating RFC 6550

This specification introduces two new RPL Control Messages to enable a RPL Aware Node (RAN) to request the establishment of a Track from self to a Target. The RAN makes its request by sending a new P-DAO Request (PDR) Message to the Root. The Root confirms with a new PDR-ACK message back to the requester RAN, see [Section 5.1](#) for more.

Section 6.7 of [\[RPL\]](#) specifies the RPL Control Message Options (CMO) to be placed in RPL messages such as the Destination Advertisement Object (DAO) message. The RPL Target Option (RTO) and the Transit Information Option (TIO) are such options.

In Non-Storing Mode, the TIO option is used in the DAO message to inform the root of the parent-child relationships within the DODAG, and the Root has a full knowledge of the DODAG structure. The TIO applies to the RTOs that precede it immediately in the message. Options may be factorized; multiple TIOs may be present to indicate multiple routes to the one or more contiguous addresses indicated in the RTOs that immediately precede the TIOs in the RPL message.

This specification introduces two new CMOs referred to as Route Projection Options (RPO) to install Projected Routes. One RPO is the Via Information Option (VIO) and the other is the Source-Routed VIO (SRVIO). The VIO installs a route on each hop along a Projected Route (in a fashion analogous to RPL Storing Mode) whereas the SRVIO installs a source-routing state at the ingress node, which uses that state to encapsulate a packet with an IPv6 Routing Header in a fashion similar to RPL Non-Storing Mode.

Like the TIO, the RPOs MUST be preceded by exactly one RTO to which they apply, and SRVIOs MAY be factorized, though VIOs MUST NOT be. Factorized contiguous SRVIOs indicate alternate paths to the Target, more in [Section 5.3](#).

This specification also introduces a new CMO to enable a RAN to advertise a selection of its candidate neighbors as siblings to the Root, using a new Sibling Information Option (SIO) as specified in [Section 5.4](#).

4. Identifying a Path

It must be noted that RPL has a concept of Instance to represent different routing topologies 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 within one routing topology.

This draft conforms the Instance model as follows:

*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 as long as it does not create a loop. A Projected Route is always preferred over a route that is learned via RPL.

*A PCE that installs a more specific (say, Traffic Engineered) and possibly complex path (i.e., a Track) towards a particular Target MUST use a Local RPL Instance (see section 5 of [RPL]) associated to that Target to identify the path.

We refer to that Local RPLInstanceID as TrackID. A TrackID MUST be unique for a particular Target address. A Track is uniquely identified within the RPL domain by the tuple (Target address, TrackID).

When packet is placed on a Track, a RPL Packet Information (RPI) is added with the TrackID as RPLInstanceID. The RPLInstanceID has the 'D' flag set, indicating that the destination address in the IPv6 header is the Target that is used to identify the Track.

*A packet that is routed over the RPL Instance associated to a Track MUST NOT be placed over a different RPL Instance again. Conversely, a packet that is placed on a Global Instance MAY be injected in a Local Instance based on a network policy and the Local Instance configuration.

A Projected Route is a serial path that may represent the end-to-end route or only a segment in a complex Track, in which case multiple Projected Routes are installed with the same tuple (Target address, TrackID) and a different segment ID each.

All properties of a Track operations are inherited from the main instance that is used to install the Track. For instance, the use of compression per [RFC8138] is determined by whether it is used in the main instance, e.g., by setting the "T" flag [TURN-ON_RFC8138] in the RPL configuration option.

5. New RPL Control Messages and Options

5.1. New P-DAO Request Control Message

The P-DAO Request (PDR) message is sent to the Root to request a new that the PCE establishes a new a projected route as a full path of a collection of segments in a Track. Exactly one Target Options MUST be present.

The format of PDR Base Object is as follows:




Figure 1: New P-DAO Request Format

TrackID: 8-bit field indicating the RPLInstanceID associated with the Track. It is set to zero upon the first request for a new

Track and then to the TrackID once the Track was created, to either renew it or destroy it.

K: The 'K' flag is set to indicate that the recipient is expected to send a PDR-ACK back.

R: The 'R' flag is set to indicate that the Requested path should be redundant.

PDRLifetime: 8-bit unsigned integer.

The requested lifetime for the Track expressed in Lifetime Units (obtained from the DODAG Configuration option).

A PDR with a fresher PDRSequence refreshes the lifetime, and a PDRLifetime of 0 indicates that the track should be destroyed.

PDRSequence: 8-bit wrapping sequence number, obeying the operation in section 7.2 of [\[RPL\]](#).

The PDRSequence is used to correlate a PDR-ACK message with the PDR message that triggered it. It is incremented at each PDR message and echoed in the PDR-ACK by the Root.

5.2. New PDR-ACK Control Message

The new PDR-ACK is sent as a response to a PDR message with the 'K' flag set. Its format is as follows:

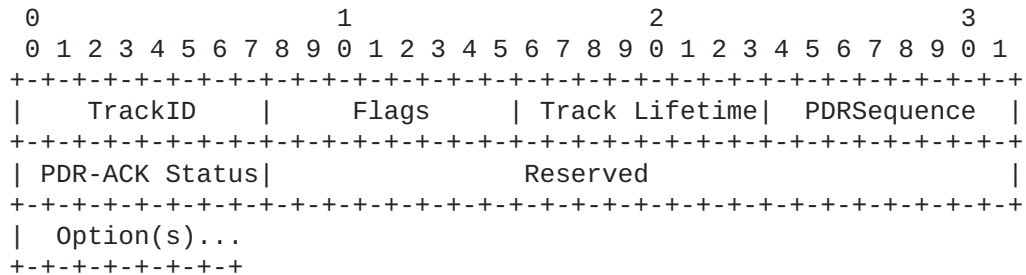


Figure 2: New PDR-ACK Control Message Format

TrackID: The RPLInstanceID of the Track that was created. Set to 0 when no Track is created.

Track Lifetime: Indicates that remaining Lifetime for the Track, 0 if the Track was destroyed or not created.

PDRSequence: 8-bit wrapping sequence number. It is incremented at each PDR message and echoed in the PDR-ACK.

PDR-ACK Status: 8-bit field indicating the completion.

The PDR-ACK Status is further substructured as indicated in [Figure 3](#).




Figure 3: PDR-ACK status Format

E: 1-bit flag. Set to indicate a rejection. When not set, a value of 0 indicates Success/Unqualified acceptance and other values indicate "not an outright rejection".

R: 1-bit flag. Reserved, MUST be set to 0 by the sender and ignored by the receiver.

Status Value: 6-bit unsigned integer. Values depending on the setting of the 'E' flag as indicated respectively in [Table 4](#) and [Table 5](#).

5.3. Route Projection Options

The RPOs indicate a series of IPv6 addresses that can be compressed using the method defined in the ["6LoWPAN Routing Header" \[RFC8138\]](#) specification using the address of the Root found in the DODAGID field of DIO messages as Compression Reference.

An RPO indicates a Projected Route that can be a serial Track in full or a segment of a more complex Track. In Non-Storing Mode, multiple RPO may be placed after a same Target Option to reflect different segments originated at this node. The Track is identified by a TrackID that is a Local RPLInstanceID to the Target of the Track.

The format of RPOs is as follows:

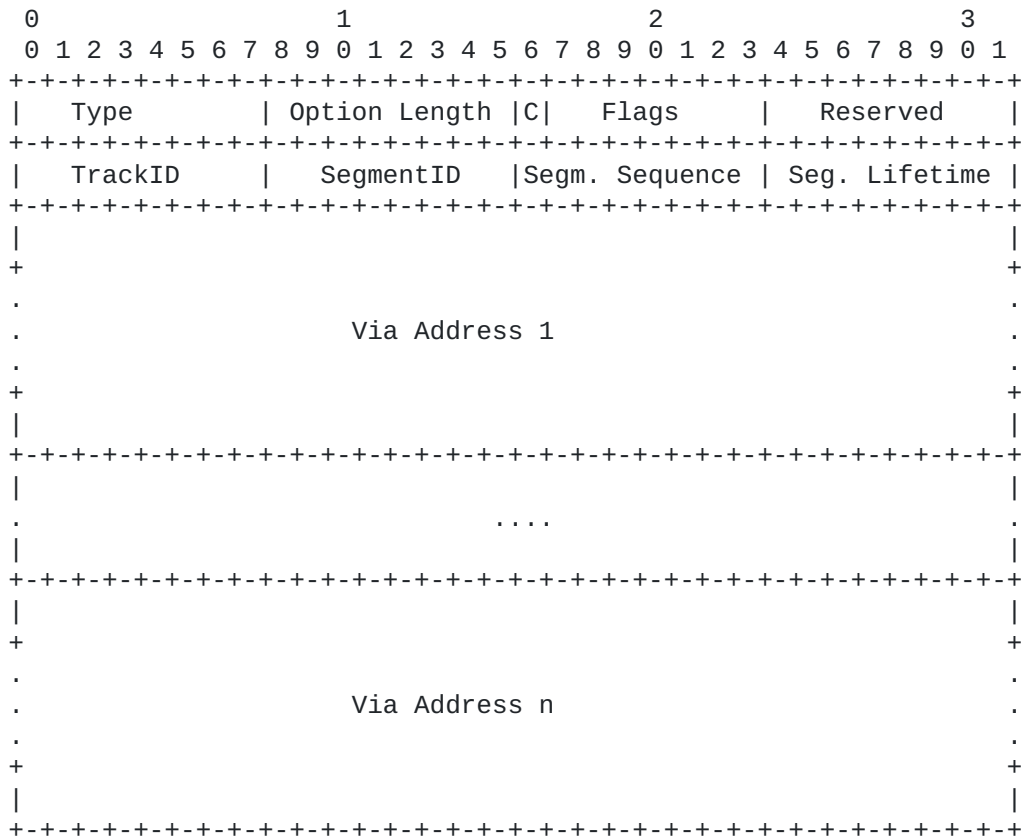


Figure 4: Via Information option format (uncompressed form)

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.

C: 1-bit flag. Set to indicate that the following Via Addresses are expressed as one or more SRH-6LoRH as defined in section 5.1 of [RFC8138]. Figure 4 illustrates the case where the "C" flag is not set, meaning that the Via Addresses are expressed in full 128 bits.

TrackID: 8-bit field indicating the topology Instance associated with the Track. The tuple (Target Address, TrackID) forms a unique ID of the Track in the RPL domain.

Track Sequence: 8-bit unsigned integer. The Track Sequence obeys the operation in section 7.2 of [RPL] and the lollipop starts at 255. The Track Sequence is set by the Root and incremented each time the Root refreshes that Track globally. A Track Sequence that is fresher than the current one deprecates any state for the Track. A Track Sequence that is current adds to any state that was learned for that Track Sequence. A RT0 with a Track Sequence that is not as fresh as the current one is ignored.

Track Lifetime: 8-bit unsigned integer. The length of time in Lifetime Units (obtained from the Configuration option) that the Track is usable. The period starts when a new Track 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.

SegmentID: 8-bit field that identifies a segment within a Track. A Value of 0 is used to signal a serial Track.

Segment Sequence: 8-bit unsigned integer. The Segment Sequence obeys the operation in section 7.2 of [RPL] and the lollipop starts at 255. When the Root of the DODAG needs to update a single segment in a Track, but does not need to modify the rest of the Track, it increments the Segment Sequence but not the Track Sequence. The segment information indicated in the RTO deprecates any state for the segment indicated by the SegmentID within the indicated Track and sets up the new information. A RTO with a Segment Sequence that is not as fresh as the current one is ignored. a RTO for a given target with the same (TrackID, Track Sequence, SegmentID, Segment Sequence) indicates a retry; it MUST NOT change the segment and MUST be propagated or answered as the first copy.

Via Address: The collection of Via Addresses along one segment, indicated in the order of the path from the ingress to the egress nodes. If the "C" flag is set, the fields Via Address 1 .. Via Address n in [Figure 4](#) are replaced by one or more of the headers illustrated in Fig. 6 of [RFC8138]. More than one SRH-6LoRH is needed if the compression of the addresses change inside the segment and different SRH-6LoRH Types are used.

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.

5.4. Sibling Information Option

The Sibling Information Option (SIO) provides indication on siblings that could be used by the Root to form Projected Routes. The format of SIOs is as follows:

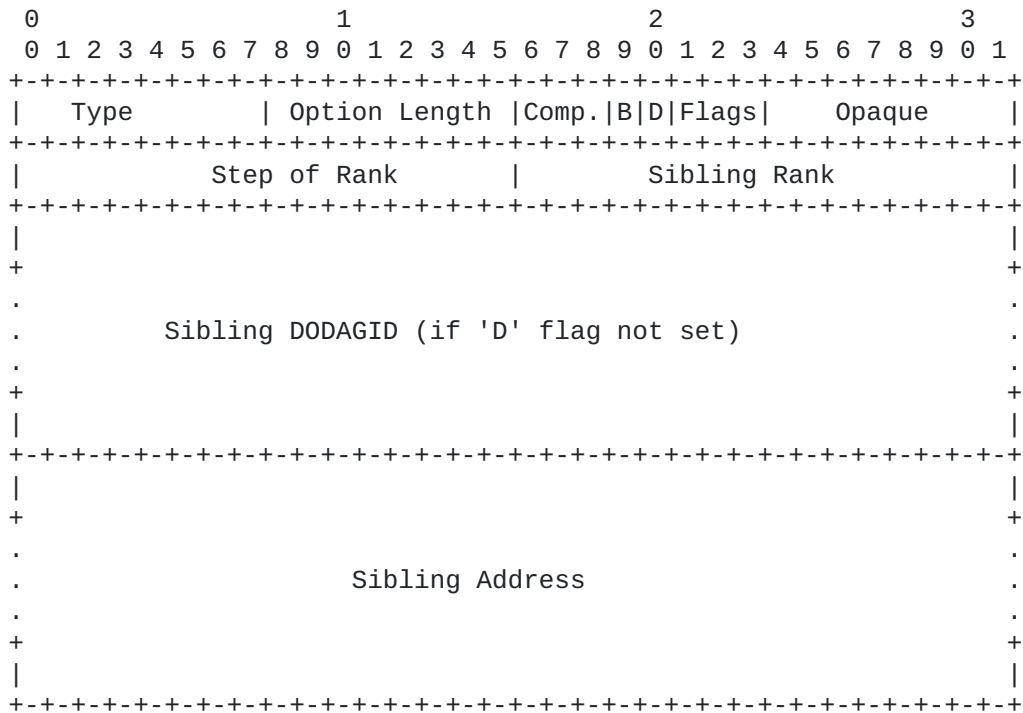


Figure 5: Sibling Information Option Format

Option Type: 0x0C (to be confirmed by IANA)

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

Compression Type: 3-bit unsigned integer. This is the SRH-6LoRH Type as defined in figure 7 in section 5.1 of [RFC8138] that corresponds to the compression used for the Sibling Address.

Reserved for Flags: MUST be set to zero by the sender and MUST be ignored by the receiver.

B: 1-bit flag that is set to indicate that the connectivity to the sibling is bidirectional and roughly symmetrical. In that case, only one of the siblings may report the SIO for the hop. If 'B' is not set then the SIO only indicates connectivity from the sibling to this node, and does not provide information on the hop from this node to the sibling.

D: 1-bit flag that is set to indicate that sibling belongs to the same DODAG. When not set, the Sibling DODAGID is indicated.

Opaque: MAY be used to carry information that the node and the Root understand, e.g., a particular representation of the Link properties such as a proprietary Link Quality Information for packets received from the sibling. An industrial Alliance that uses RPL for a particular use / environment MAY redefine the use of this field to fit its needs.

Step of Rank: 16-bit unsigned integer. This is the Step of Rank [RPL] as computed by the Objective Function between this node and the sibling.

Sibling Rank:

16-bit unsigned integer. When non-zero, this is the Rank [[RPL](#)] as exposed by the sibling in DIO messages.

Sibling DODAGID: 2 to 16 bytes, the DODAGID of the sibling in a [[RFC8138](#)] compressed form as indicated by the Compression Type field. This field is present when the 'D' flag is not set.

Sibling Address: 2 to 16 bytes, the IPv6 Address of the sibling in a [[RFC8138](#)] compressed form as indicated by the Compression Type field.

An SIO MAY be immediately followed by a DAG Metric Container. In that case the DAG Metric Container provides additional metrics for the hop from the Sibling to this node.

6. Projected DAO

This draft adds a capability to RPL whereby the Root of a DODAG projects a route by sending one or more extended DAO message called Projected-DAO (P-DAO) messages 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 RPL Target Option(s) (RTO) 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 exactly one RTO and either one VIO or one or more SRVIOs following it. There can be at most one such sequence of RTOs and then RPOs.

Like a classical DAO message, a P-DAO causes a change of state only if it is "new" per section 9.2.2. "Generation of DAO Messages" of the [RPL specification](#) [[RPL](#)]; this is determined using the Track Sequence and the Segment Sequence information from the RPO as opposed to the Path Sequence from 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 Projected Routes, the Storing Mode and the Non-Storing Mode.

*The Non-Storing Mode is discussed in [Section 6.1](#). 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.

*The Storing Mode is discussed in [Section 6.2](#). 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

next source routed hop and fails to locate it as a neighbor or a neighbor of a neighbor, then it MUST ensure that it has another Projected Route to the next loose hop under the control of the same route computation system, otherwise the P-DAO is rejected.

When forwarding a packet to a destination for which the router determines that routing happens via the Target, the router inserts the source routing header in the packet to reach the Target. 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 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 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 [[RFC8138](#)]. 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.

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 [[RPL](#)]. 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.

6.2. Storing-Mode Projected Route

As illustrated in [Figure 7](#), the Storing Mode route projection is 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.

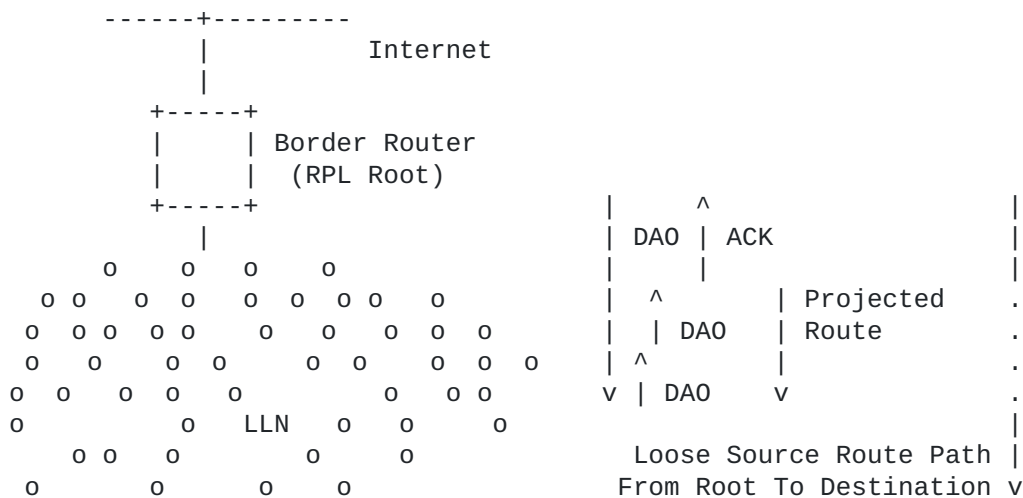


Figure 7: Projecting a route

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 Projected 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 Projected 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 Projected Route that is broken.

7. Security Considerations

This draft uses messages that are already present in RPL [RPL] 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?

8. IANA Considerations

8.1. New RPL Control Codes

This document extends the IANA Subregistry created by RFC 6550 for RPL Control Codes as indicated in [Table 1](#):

Code	Description	Reference
0x09	Projected DAO Request (PDR)	This document
0x0A	PDR-ACK	This document

Table 1: New RPL Control Codes

8.2. New RPL Control Message Options

This document extends the IANA Subregistry created by RFC 6550 for RPL Control Message Options as indicated in [Table 2](#):

Value	Meaning	Reference
0x0B	Via Information option	This document
0x0C	Source-Routed Via Information option	This document
0x0D	Sibling Information option	This document

Table 2: RPL Control Message Options

8.3. New SubRegistry for the Projected DAO Request (PDR) Flags

IANA is required to create a registry for the 8-bit Projected DAO Request (PDR) Flags field. Each bit is tracked with the following qualities:

*Bit number (counting from bit 0 as the most significant bit)

*Capability description

*Reference

Registration procedure is "Standards Action" [[RFC8126](#)]. The initial allocation is as indicated in [Table 3](#):

Bit number	Capability description	Reference
0	PDR-ACK request (K)	This document
1	Requested path should be redundant (R)	This document

Table 3: Initial PDR Flags

8.4. New SubRegistry for the PDR-ACK Flags

IANA is required to create an subregistry for the 8-bit PDR-ACK Flags field. Each bit is tracked with the following qualities:

*Bit number (counting from bit 0 as the most significant bit)

*Capability description

*Reference

Registration procedure is "Standards Action" [[RFC8126](#)]. No bit is currently defined for the PDR-ACK Flags.

8.5. New Subregistry for the PDR-ACK Acceptance Status values

IANA is requested to create a new subregistry for the PDR-ACK Acceptance Status values.

*Possible values are 6-bit unsigned integers (0..63).

*Registration procedure is "Standards Action" [[RFC8126](#)].

*Initial allocation is as indicated in [Table 4](#):

Value	Meaning	Reference
0	Unqualified acceptance	This document

Table 4: Acceptance values of the PDR-ACK Status

8.6. New Subregistry for the PDR-ACK Rejection Status values

IANA is requested to create a new subregistry for the PDR-ACK Rejection Status values.

*Possible values are 6-bit unsigned integers (0..63).

*Registration procedure is "Standards Action" [[RFC8126](#)].

*Initial allocation is as indicated in [Table 5](#):

Value	Meaning	Reference
0	Unqualified rejection	This document

Table 5: Rejection values of the PDR-ACK Status

8.7. New SubRegistry for the Route Projection Options (RPO) Flags

IANA is requested to create a new subregistry for the 5-bit Route Projection Options (RPO) Flags field. Each bit is tracked with the following qualities:

*Bit number (counting from bit 0 as the most significant bit)

*Capability description

*Reference

Registration procedure is "Standards Action" [[RFC8126](#)]. No bit is currently defined for the Route Projection Options (RPO) Flags.

8.8. New SubRegistry for the Sibling Information Option (SIO) Flags

IANA is required to create a registry for the 5-bit Sibling Information Option (SIO) Flags field. Each bit is tracked with the following qualities:

*Bit number (counting from bit 0 as the most significant bit)

*Capability description

*Reference

Registration procedure is "Standards Action" [[RFC8126](#)]. The initial allocation is as indicated in [Table 6](#):

Bit number	Capability description	Reference
0	Connectivity is bidirectional (B)	This document

Table 6: Initial SIO Flags

8.9. Error in Projected Route ICMPv6 Code

In some cases RPL will return an ICMPv6 error message when a message cannot be forwarded along a Projected 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.

9. Acknowledgments

The authors wish to acknowledge JP Vasseur, Remy Liubing, James Pylakutty and Patrick Wetterwald for their contributions to the ideas developed here.

10. 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>>.
- [RPL] 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>>.
- [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>>.
- [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>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

11. Informative References

- [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>>.
- [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>>.
- [6TISCH-ARCHI] Thubert, P., "An Architecture for IPv6 over the TSCH mode of IEEE 802.15.4", Work in Progress, Internet-Draft, draft-ietf-6tisch-architecture-29, August 27, 2020, <<https://tools.ietf.org/html/draft-ietf-6tisch-architecture-29>>.
- [RAW-ARCHI] Thubert, P., Papadopoulos, G., and R. Buddenberg, "Reliable and Available Wireless Architecture/Framework", Work in Progress, Internet-Draft, draft-pthubert-raw-architecture-04, July 6, 2020, <<https://tools.ietf.org/html/draft-pthubert-raw-architecture-04>>.
- [TURN-ON_RFC8138] Thubert, P. and L. Zhao, "Configuration option for RFC 8138", Work in Progress, Internet-Draft, draft-thubert-roll-turnon-rfc8138-03, July 8, 2019, <<https://tools.ietf.org/html/draft-thubert-roll-turnon-rfc8138-03>>.
- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.
- [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>>.
- [PCE] IETF, "Path Computation Element", , <<https://datatracker.ietf.org/doc/charter-ietf-pce/>>.

Appendix A. Applications

A.1. Loose Source Routing in Non-storing Mode

A RPL implementation operating in a very constrained LLN typically uses the Non-Storing Mode of Operation as represented in [Figure 8](#). In that mode, a RPL node indicates a parent-child relationship to the Root, using a Destination Advertisement Object (DAO) that is

unicast from the node directly to the Root, and the Root typically builds a source routed path to a destination down the DODAG by recursively concatenating this information.

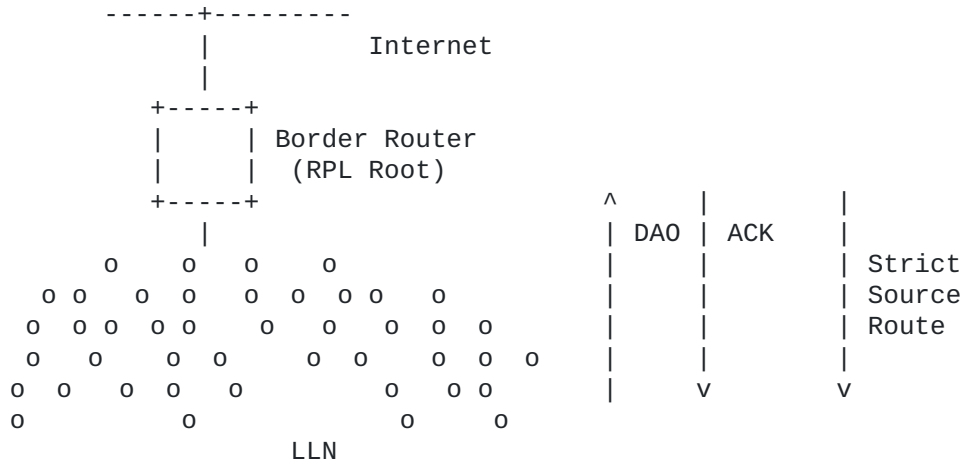


Figure 8: RPL non-storing mode of operation

Based on the parent-children relationships expressed in the non-storing DAO messages, the Root possesses topological information about the whole network, though this information is limited to the structure of the DODAG for which it is the destination. A packet that is generated within the domain will always reach the Root, which can then apply a source routing information to reach the destination if the destination is also in the DODAG. Similarly, a packet coming from the outside of the domain for a destination that is expected to be in a RPL domain reaches the Root.

It results that the Root, or then some associated centralized computation engine such as a PCE, can determine the amount of packets that reach a destination in the RPL domain, and thus the amount of energy and bandwidth that is wasted for transmission, between itself and the destination, as well as the risk of fragmentation, any potential delays because of a paths longer than necessary (shorter paths exist that would not traverse the Root).

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 9](#):

*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.

*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 9: 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 10: 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 11](#), if the source is node 41 and the destination is node 52, then the common parent is node 22.




Figure 11: 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 12: 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 of 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 13: P-DAO-ACK to Root

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




Figure 14: Projected Transversal Route

Authors' Addresses

Pascal Thubert (editor)
 Cisco Systems, Inc
 Building D
 45 Allee des Ormes - BP1200
 06254 Mougins - Sophia Antipolis
 France

Phone: [+33 497 23 26 34](tel:+33497232634)
 Email: pthubert@cisco.com

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

Phone: [+91-080-49160700](tel:+9108049160700)
 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](tel:+18006355461)
 Email: matthew.gillmore@itron.com