ROLL
Internet-Draft
Intended status: Standards Track
Expires: 31 May 2021

P. Thubert, Ed.
Cisco Systems
R.A. Jadhav
Huawei Tech
M. Gillmore
Itron

27 November 2020

Root initiated routing state in RPL
draft-ietf-roll-dao-projection-15

Abstract

This document extends RFC 6550 and RFC 6553 to enable 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 Routing Header or in terms of path length, which impacts both the latency and the packet delivery ratio.

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 31 May 2021.

Copyright Notice

Table of Contents

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. 6TiSCH uses RPL for its routing operations.

The "6TiSCH Architecture" [6TiSCH-ARCHI] 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, the PCE with a global visibility on the system can compute direct Peer to Peer (P2P) routes that are optimized for the needs expressed by an objective function. This document specifies protocol extensions to RPL [RPL] that enable the Root of a main DODAG to install centrally-computed routes inside the DODAG on behalf of a PCE.

This specification expects that the main RPL Instance is operated in RPL Non-Storing Mode of Operation (MOP) 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.

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. Projected Routes can be used to reduce the size of the source routing headers with loose source routing operations down the main RPL DODAG. Projected Routes can also be used to build transversal routes for route optimization and Traffic Engineering purposes, between nodes of the DODAG.

A Projected Route may be installed in either Storing and Non-Storing Mode, potentially resulting in hybrid situations where the Mode of the Projected Route is different from that of the main RPL Instance. A Projected Route may be a stand-alone end-to-end path 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 potentially complex path with redundant forwarding solutions along the way. With this specification, a Track is a DODAG formed by a RPL local Instance that is rooted at the Track Ingress. If there is a single Track Egress, then the Track is reversible to form another DODAG by reversing the direction of each edge. A node at the ingress of more than one Segment in a Track may use one or more of these Segments to forward a packet inside the Track.

The "Reliable and Available Wireless (RAW) Architecture/Framework" [RAW-ARCHI] defines the Path Selection Engine (PSE) that adapts the use of the path redundancy within a Track to defeat the diverse causes of packet loss.

The PSE is a dataplane extension of the PCE; it controls the forwarding operation of the packets within a Track, using Packet ARQ, Replication, Elimination, and Overhearing (PAREO) functions over the Track segments, to provide a dynamic balance between the reliability and availability requirements of the flows and the need to conserve energy and spectrum.

The time scale at which the PCE (re)computes the Track can be long, using long-term statistical metrics to perform global optimizations at the scale of the whole network. Conversely, the PSE makes forwarding decisions at the time scale of one or a small collection of packets, based on a knowledge that is limited in scope to the Track itself, so it can be refreshed at a fast pace.

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 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. Glossary

This document often uses the following acronyms:

CMO: Control Message Option
DAO: Destination Advertisement Object
DAG: Directed Acyclic Graph
DODAG: Destination-Oriented Directed Acyclic Graph; A DAG with only one vertex (i.e., node) that has no outgoing edge (i.e., link)
LLN: Low-Power and Lossy Network
NMPR: Non-Storing Mode Projected Route
MOP: RPL Mode of Operation
P-DAO: Projected DAO
PDR: P-DAO Request
RAN: RPL-Aware Node (either a RPL Router or a RPL-Aware Leaf)
RAL: RPL-Aware Leaf
RH: Routing Header
RPI: RPL Packet Information
RTO: RPL Target Option
RUL: RPL-Unaware Leaf
SIO: RPL Sibling Information Option
SR-VIO: A Source-Routed Via Information Option, used in Non-Storing Mode P-DAO messages.
SMPR: Storing Mode Projected Route
TIO: RPL Transit Information Option
SF-VIO: A Via Information Option, used in Storing Mode P-DAO messages.
VIO: A Via Information Option; it can be a SF-VIO or an SR-VIO.

2.3. Other Terms

Projected Route: A RPL Projected Route is a RPL route that is computed remotely by a PCE, and installed and maintained by a RPL Root on behalf of the PCE.

Projected DAO: A DAO message used to install a Projected Route.

Track: A DODAG that provides a complex path from or to a Root that is the destination of the DODAG. The Root is the Track Ingress, and the forward direction for packets is down the DODAG, from the Track Ingress to one of the possibly multiple Track Egress Nodes.

TrackID: A RPL Local InstanceID with the 'D' bit set to 0. The TrackID is associated with the IPv6 Address of the Track Ingress that is used to signal the DODAG Root.

2.4. 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].

3. Extending RFC 6550

3.1. Projected DAO

Section 6 of [RPL] introduces the RPL Control Message Options (CMO), including the RPL Target Option (RTO) and Transit Information Option (TIO), which can be placed in RPL messages such as the Destination Advertisement Object (DAO). This specification extends the DAO message with the Projected DAO (P-DAO); a P-DAO message signals a Projected Route to one or more Targets using the new CMOs presented therein. This specification enables to combine one or more Projected Routes into a DODAG called a Track, that is traversed to reach the Targets.

The Track is assimilated with the DODAG formed for a Local RPL Instance. The local RPLInstanceID of the Track is called the TrackID, more in Section 7.2. A P-DAO message for a Track signals the TrackID in the RPLInstanceID field. The Track Ingress is signaled in the DODAGID field of the Projected DAO Base Object; that field is elided in the case of the main RPL Instance. The Track Ingress is the Root of the Track, as shown in Figure 1. .

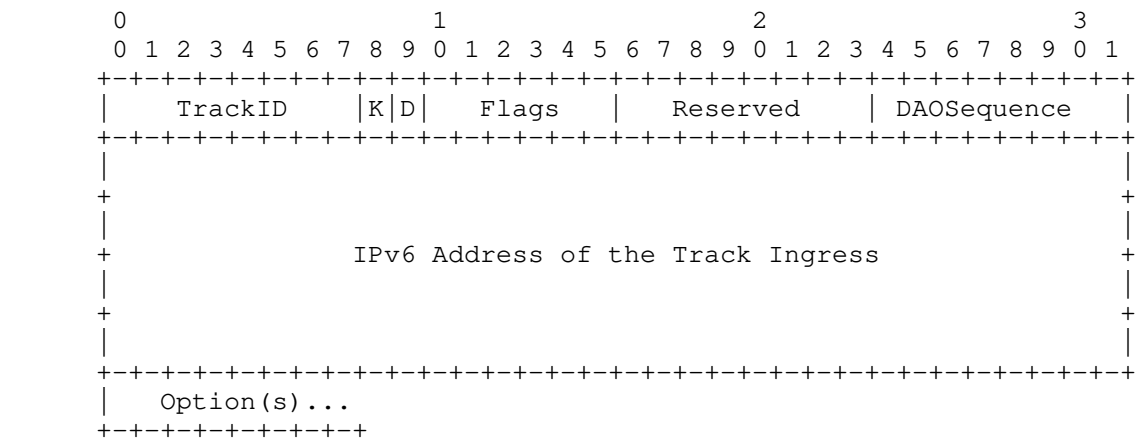

Figure 1: Projected DAO Format for a Track

In RPL Non-Storing Mode, the TIO and RTO are combined in a DAO message to inform the DODAG Root of all the edges in the DODAG, which are formed by the directed parent-child relationships. Options may be factorized; multiple RTOs may be present to signal a collection of children that can be reached via the parent(s) indicated in the TIO(s) that follows the RTOs. This specification generalizes the case of a parent that can be used to reach a child with that of a whole Track through which both children and siblings of the Track Egress are reachable.

New CMOs called the Via Information Options (VIO) are introduced for use in P-DAO messages as a multihop alternative to the TIO. One VIO is the Stateful Via Information Option (SF-VIO); the SF-VIO installs Storing Mode Projected Route (SMPR) along a strict segment. The other is the Source-Routed SF-VIO (SR-VIO); the SR-VIO installs a Non-Storing Mode Projected Route (NMPR) at the Track Ingress, which uses that state to encapsulate a packet with a Routing Header (RH) to the Track Egress.

Like in a DAO message, the RTOs can be factorized in a P-DAO, but the Via Options cannot. A P-DAO contains one or more RTOs that indicate the destinations that can be reached via the Track, and exactly one Via Option that signals a sequence of nodes. In Non-Storing Mode, the Root sends the P-DAO to the Track Ingress where the source-routing state is stored. In Storing Mode, the P-DAO is sent to the Track Egress and forwarded along the Segment in the reverse direction, installing a Storing Mode state to the Track Egress at each hop. In both cases the Track Ingress is the owner of the Track, and it generates the P-DAO-ACK when the installation is successful.

3.2. Sibling Information Option

This specification adds another CMO called the Sibling Information Option (SIO) that is used by a RPL Aware Node (RAN) to advertise a selection of its candidate neighbors as siblings to the Root, more in Section 6.4. The sibling selection process is out of scope.

3.3. P-DAO Request

Two new RPL Control Messages are also introduced, to enable a RAN to request the establishment of a Track between self as the Track Ingress Node and a Track Egress. 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 6.1 for more. A positive PDR-ACK indicates that the Track was built and that the Roots commits to maintain the Track for the negotiated lifetime. In the case of a complex Track, each Segment is maintained independently and asynchronously by the Root, with its own lifetime that may be shorter, the same, or longer than that of the Track. The Root may use an asynchronous PDR-ACK with an negative status to indicate that the Track was terminated before its time.

3.4. Extending the RPI

Sending a Packet within a RPL Local Instance requires the presence of the abstract RPL Packet Information (RPI) described in section 11.2. of [RPL] in the outer IPv6 Header chain (see [USEofRPLinfo]). The RPI carries a local RPLInstanceID which, in association with either the source or the destination address in the IPv6 Header, indicates the RPL Instance that the packet follows.

This specification extends [RPL] to create a new flag that signals that a packet is forwarded along a projected route.

Projected-Route 'P': 1-bit flag. It is set to 1 if this packet is sent over a projected route and set to 0 otherwise.

4. Extending RFC 6553

"The RPL Option for Carrying RPL Information in Data-Plane Datagrams" [RFC6553]describes the RPL Option for use among RPL routers to include the abstract RPL Packet Information (RPI) described in section 11.2. of [RPL] in data packets.

The RPL Option is commonly referred to as the RPI though the RPI is really the abstract information that is transported in the RPL Option. [USEofRPLinfo] updated the Option Type from 0x63 to 0x23.

This specification modifies the RPL Option to encode the 'P' flag as follows:

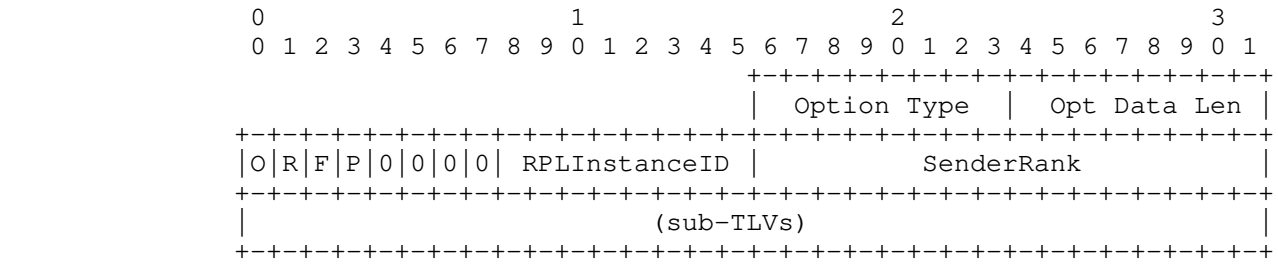

Figure 2: Extended RPL Option Format

Option Type: 0x23 or 0x63, see [USEofRPLinfo]

Opt Data Len: See [RFC6553]

'O', 'R' and 'F' flags: See [RFC6553]. Those flags MUST be set to 0 by the sender and ignored by the receiver if the 'P' flag is set.

Projected-Route 'P': 1-bit flag as defined in Section 3.4.

RPLInstanceID: See [RFC6553]. Indicates the TrackId if the 'P' flag is set.

SenderRank: See [RFC6553]. This field MUST be set to 0 by the sender and ignored by the receiver if the 'P'flag is set.

5. Extending RFC 8138

Section 6.3 of [RFC8138] presents the formats of the 6LoWPAN Routing Header of type 5 (RPI-6LoRH) that compresses the RPI for normal RPL operation. The format of the RPI-6LoRH is not suited for Projected routes since the O,R,F flags are not used and the Rank is unknown and ignored.

This specification introduces a new 6LoRH, the P-RPI-6LoRH, with a type of 7. The P-RPI-6LoRH header is usually a a Critical 6LoWPAN Routing Header, but it can be elective as well if an SRH-6LoRH is present and controls the routing decision.

The P-RPI-6LoRH is designed to compress the RPI along RPL Projected Routes. It sformat is as follows:

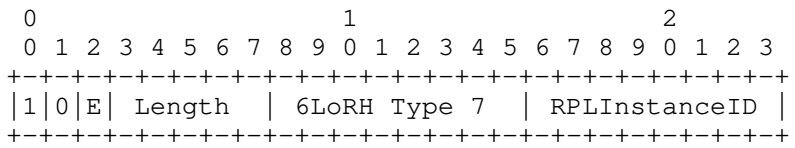

Figure 3: P-RPI-6LoRH Format

Elective 'E': See [RFC8138]. The 'E' flag is set to 1 to indicate an Elective 6LoRH, meaning that it can be ignored when forwarding.

6. New RPL Control Messages and Options

6.1. New P-DAO Request Control Message

The P-DAO Request (PDR) message is sent by a Node in the main DODAG to the Root. It is a request to establish or refresh a Track. Exactly one RTO MUST be present in a PDR. The RTO signals the Track Egress, more in Section 7.1.

The RPL Control Code for the PDR is 0x09, to be confirmed by IANA. The format of PDR Base Object is as follows:

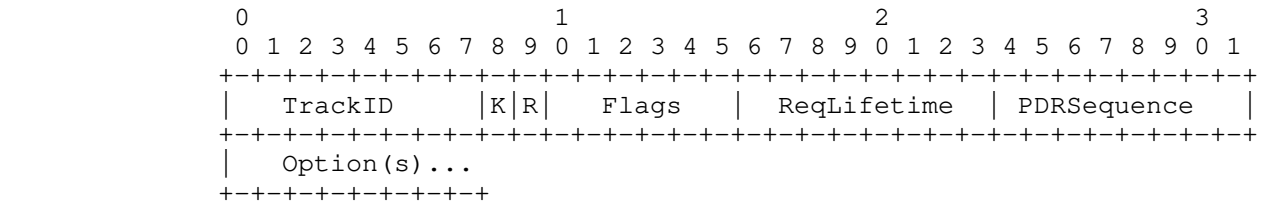

Figure 4: New P-DAO Request Format

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 of 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 request a Complex Track for redundancy.

Flags: Reserved. The Flags field MUST initialized to zero by the sender and MUST be ignored by the receiver

ReqLifetime: 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.

6.2. New PDR-ACK Control Message

The new PDR-ACK is sent as a response to a PDR message with the 'K' flag set. The RPL Control Code for the PDR-ACK is 0x0A, to be confirmed by IANA. Its format is as follows:

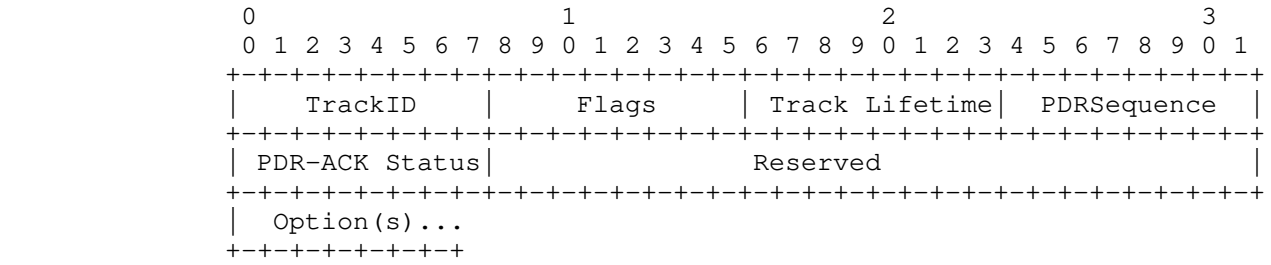

Figure 5: New PDR-ACK Control Message Format

TrackID: The RPLInstanceID of the Track that was created. The value of 0x00 is used to when no Track was created.

Flags: Reserved. The Flags field MUST initialized to zero by the sender and MUST be ignored by the receiver

Track Lifetime: Indicates that remaining Lifetime for the Track, expressed in Lifetime Units; the value of zero (0x00) indicates that 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 substructured as indicated in Figure 6:

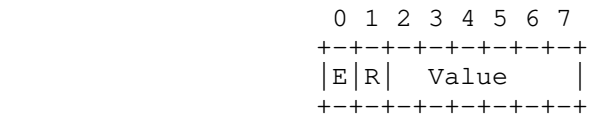

Figure 6: PDR-ACK status Format

E: 1-bit flag. Set to indicate a rejection. When not set, the 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, see Table 7 and Table 8.

Reserved: The Reserved field MUST initialized to zero by the sender and MUST be ignored by the receiver

6.3. Via Information Options

An Via Option signals the ordered list of IPv6 Via Addresses that constitutes the hops of either a Serial Track or a Segment of a more Complex Track. An Via Option MUST contain at least one Via Address, and a Via Address MUST NOT be present more than once, otherwise the Via Option MUST be ignored. The format of the Via Options is as follows:

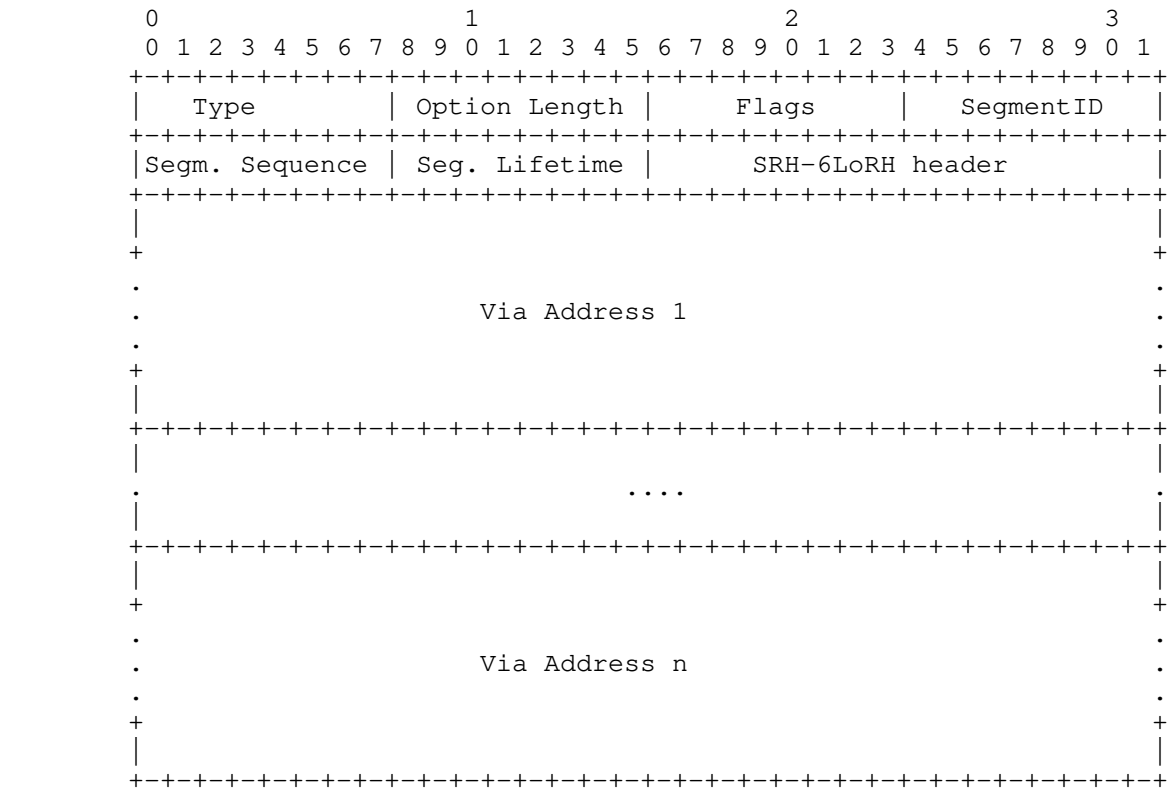

Figure 7: Via Information Option format (uncompressed form)

Option Type: 0x0B for SF-VIO, 0x0C for SR-VIO (to be confirmed by IANA)

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

SegmentID: 8-bit field that identifies a Segment within a Track or the main DODAG as indicated by the TrackID field. The value of 0 is used to signal a Serial Track, i.e., made of a single segment.

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 refresh or update a Segment in a Track, it increments the Segment Sequence individually for that Segment.

The Segment information indicated in the Via Option deprecates any state for the Segment indicated by the SegmentID within the indicated Track and sets up the new information.

An Via Option with a Segment Sequence that is not as fresh as the current one is ignored.

A VIO for a given DODAGID with the same (TrackID, SegmentID, Segment Sequence) indicates a retry; it MUST NOT change the Segment and MUST be propagated or answered as the first copy.

Segment Lifetime: 8-bit unsigned integer. The length of time in Lifetime Units (obtained from the Configuration option) that the Segment is usable.

The period starts when a new Segment Sequence is seen. The value of 255 (0xFF) represents infinity. The value of zero (0x00) indicates a loss of reachability.

A P-DAO message that contains a Via Information option with a Segment Lifetime of zero is referred as a No-Path P-DAO in this document.

SRH-6LoRH header: The first 2 bytes of the (first) SRH-6LoRH as shown in Figure 6 of [RFC8138]. A 6LoRH Type of 4 means that the VIA Addresses are provided in full with no compression.

Via Address: An IPv6 addresse along the Segment.

In a SF-VIO, the list is a strict path between direct neighbors, from the segment ingress to egress, both included. In an SR-VIO, the list starts at the first hop and ends at a Track Egress. The list in an SR-VIO may be loose, provided that each listed node has a path to the next listed node, e.g., via a segment or another Track.

In the case of a SF-VIO, or if [RFC8138] is not used in the data packets, then the Root MUST use only one SRH-6LoRH per Via Option, and the compression is the same for all the addresses, as shown in Figure 7.

In case of an SR-VIO, and if [RFC8138] is in use in the main DODAG, then the Root SHOULD optimize the size of the SR-VIO; more than one SRH-6LoRH may be present, e.g., if the compression level changes inside the Segment and different SRH-6LoRH Types are required. The content of the SR-VIO starting at the first SRH-6LoRH header is thus verbatim the one that the Track Ingress places in the packet encapsulation to reach the Track Ingress.

6.4. Sibling Information Option

The Sibling Information Option (SIO) provides indication on siblings that could be used by the Root to form Projected Routes. One or more SIO(s) may be placed in the DAO messages that are sent to the Root in Non-Storing Mode.

The format of the SIO is as follows:

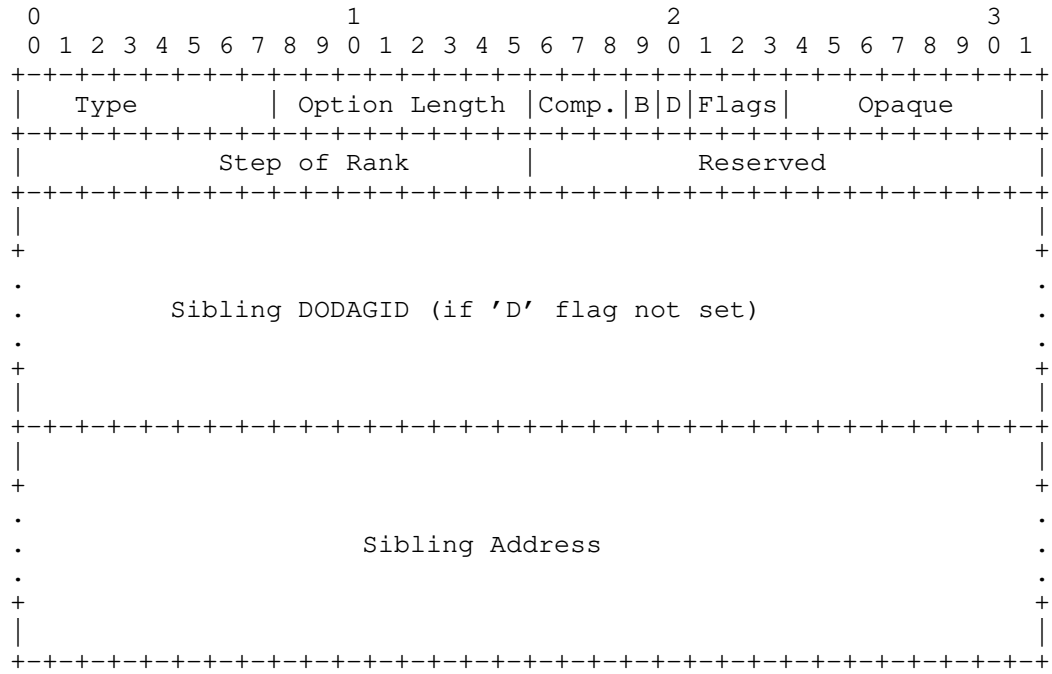

Figure 8: Sibling Information Option Format

Option Type: 0x0D (to be confirmed by IANA)

Option Length: In bytes, the size of the option.

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 and its DODAGID if resent. The Compression refernce is the Root of the main DODAG.

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.

Flags: Reserved. The Flags field MUST initialized to zero by the sender and MUST be ignored by the receiver

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.

Reserved: The Reserved field MUST initialized to zero by the sender and MUST be ignored by the receiver

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.

7. Projected DAO

This draft adds a capability to RPL whereby the Root of a main DODAG installs a Track as a collection of Projected Routes, using a Projected-DAO (P-DAO) message to maintain each individual route. The P-DAO signals a collection of Targets in the RPL Target Option(s) (RTO). Those Targets can be reached via a sequence of routers indicated in a Via Information Option (VIO). A P-DAO message MUST contain exactly one VIO, which is either a SF-VIO or an SR-VIO, and

MUST follow one or more RTOs. There can be at most one such sequence of RTO(s) and an Via Option. A track is indentified by a tupple DODAGID, TrackID and each route within a Track is indexed by a SegmentID.

A P-DAO MUST be sent from the address of the Root that serves as DODAGID for the main DODAG. It MUST be sent to a GUA or a ULA of either the ingress or the egress of the Segment, more below. If the 'K' Flag is present in the P-DAO, and unless the P-DAO does not reach it, the ingress of the Segment is the node that acknowledges the message, using a DAO-ACK that MUST be sent back to the address that serves as DODAGID for the main DODAG.

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 Segment Sequence information from the Via Option as opposed to the Path Sequence from a TIO. Also, a Segment Lifetime of 0 in an Via Option indicates that the projected route associated to the Segment 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 7.5. A Non-Storing Mode P-DAO carries an SR-VIO with the loose list of Via Addresses that forms a source-routed Segment to the Track Egress. The recipient of the P-DAO is the Track Ingress; it MUST install a source-routed state to the Track Egress and reply to the Root directly using a DAO-ACK message if requested to.

* The Storing Mode is discussed in Section 7.6. A Storing Mode P-DAO carries a SF-VIO with the strict list of Via Addresses from the ingress to the egress of the Segment in the data path order. The routers listed in the Via Addresses, except the egress, MUST install a routing state to the Target(s) via the next Via Address in the SF-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.

In case of a forwarding error along a Projected Route, an ICMP error is sent to the Root with a new Code "Error in Projected Route" (See Section 9.12). The Root can then modify or remove the Projected 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 RH if one is present, and this hop of the RH 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 ICMP message.

The sender and exact operation depend on the Mode and is described in Section 7.5 and Section 7.6 respectively.

7.1. Requesting a Track

A Node is free to ask the Root for a new Track at any time. This is done with a PDR message, that indicates in the Requested Lifetime field the duration for which the Track should be established. Upon a PDR, the Root MAY install the necessary Segments, in which case it answers with a PDR-ACK indicating the granted Track Lifetime. All the Segments MUST be of a same mode, either Storing or Non-Storing. All the Segments MUST be created with the same TrackID and the same DODAGID signaled in the P-DAO.

The Root is free to design the Track as it wishes, and to change the Segments overtime to serve the Track as needed, without notifying the resquesting Node. The Segment Lifetime in the P-DAO messages does not need to be aligned to the Requested Lifetime in the PDR, or between P-DAO messages for different Segments. The Root may use shorter lifetimes for the Segments and renew them faster than the Track is, or longer lifetimes in which case it will need to tear down the Segments if the Track is not renewed.

When the Track Lifetime that was returned in the PDR-ACK is close to elapse, the resquesting Node needs to resend a PDR using the TrackID in the PDR-ACK to extend the lifetime of the Track, else the Track will time out and the Root will tear down the whole structure.

If the Track fails and cannot be restored, the Root notifies the resquesting Node asynchronously with a PDR-ACK with a Track Lifetime of 0, indicating that the Track has failed, and a PDR-ACK Status indicating the reason of the fault.

7.2. Identifying a Track

RPL defines the concept of an Instance to signal an individual routing topology 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 leverages the RPL Instance model as follows:

* The Root MAY use P-DAO messages to add better routes in the main (Global) Instance in conformance with the routing objectives in that Instance. To achieve this, the Root MAY install an SMPR along a path down the main Non-Storing Mode DODAG. This enables a loose source routing and reduces the size of the Routing Header, see Appendix A.1.

  When adding an SMPR to the main RPL Instance, the Root MUST set the RPLInstanceID field of the P-DAO message (see section 6.4.1. of [RPL]) to the RPLInstanceID of the main DODAG, and MUST NOT use the DODAGID field. A Projected Route provides a longer match to the Target Address than the default route via the Root, so it is preferred.

  Once the Projected Route is installed, the intermediate nodes listed in the SF-VIO after first one (i.e. The ingress) can be elided from the RH in packets sent along the Segment signaled in the P-DAO. The resulting loose source routing header indicates (one of) the Target(s) as the next entry after the ingress.

* The Root MAY also use P-DAO messages to install a specific (say, Traffic Engineered) path as a Serial or as a Complex Track, to a particular endpoint that is the Track Egress. In that case, the Root MUST install a Local RPL Instance (see section 5 of [RPL]).

  In a that case, the TrackID MUST be unique for the Global Unique IPv6 Address (GUA) or Unique-Local Address (ULA) of the Track Ingress that serves as DODAGID for the Track. This way, a Track is uniquely identified by the tuple (DODAGID, TrackID) where the TrackID is always represented with the 'D' flag set to 0.

  The Track Egress Address and the TrackID MUST be signaled in the P-DAO message as shown in Figure 1.

7.3. Installing a Track

A Storing Mode P-DAO contains an SF-VIO that signals the strict sequence of consecutive nodes to form a segment between a segment ingress and a segment egress (both included). It installs a route of a higher precedence along the segment towards the Targets indicated in the Target Options. The segment is included in a DODAG indicated by the P-DAO Base Object, that may be the one formed by the main RPL Instance, or a Track associated with a local RPL Instance. A Track Egress is signaled as a Target in the P-DAO, and as the last entry is an SF-VIO of a last segment towards that Egress.

A Non-Storing Mode P-DAO signals a strict or loose sequence of nodes between the Track Ingress (excluded) and a Track Egress (included). It installs a source-routed path of a higher precedence within the Track indicated by the P-DAO Base Object, towards the Targets indicated in the Target Options. The source-routed path requires a Source-Routing header which implies an encapsulation to add the SRH to an existing packet.

The next entry in the sequence must be either a neighbor of the previous entry, or reachable as a Target via another Projected Route, either Storing or Non-Storing. If it is reachable over a Storing Mode Projected Route, the next entry in the loose sequence is the Target of a previous segment and the ingress of a next segment; the segments are associated with the same Track, which avoids the need of an encapsulation. Conversely, if it is reachable over a Non-Storing Mode Projected Route, the next loose source routed hop of the inner Track is a Target of a previous Track and the ingress of a next Track, which requires a de- and a re-encapsulation.

A Serial Track is installed by a single Projected Routes that signals the sequence of consecutive nodes, either in Storing or Non-Storing Mode. If can be a loose Non-Storing Mode Projected Route, in which case the next loose entry must recursively be reached over a Serial Track.

A Complex Track can be installed as a collection of Projected Routes with the same DODAGID and Track ID. The Ingress of a Non-Storing Mode Projected Route must be the owner of the DODAGID. The Ingress of a Storing Mode Projected Route must be either the owner of the DODAGID, or the egress of a preceding Storing Mode Projected Route in the same Track. In the latter case, the Targets of the Projected Route must be Targets of the preceding Projected Route to ensure that they are visible from the track Ingress.

7.4. Forwarding Along a Track

This draft leverages the RPL Forwarding model follows:

* In the data packets, the Track DODAGID and the TrackID MUST be respectively signaled as the IPv6 Source Address and the RPLInstanceID field of the RPI that MUST be placed in the outer chain of IPv6 Headers.

  The RPI carries a local RPLInstanceID called the TrackID, which, in association with the DODAGID, indicates the Track along which the packet is forwarded.

The 'D' flag in the RPLInstanceID MUST be set to 0 to indicate that the source address in the IPv6 header is set ot the DODAGID, more in Section 7.4.

* This draft conforms the principles of [USEofRPLinfo] with regards to packet forwarding and encapsulation along a Track.

  - In that case, the Track is the DODAG, the Track Ingress is the Root, and the Track Egress is a RAL, and neighbors of the Track Egress that can be reached via the Track are RULs. The encapsulation rules in [USEofRPLinfo] apply.

  - If the Track Ingress is the originator of the packet and the Track Egress is the destination of the packet, there is no need for an encapsulation.

  - So the Track Ingress must encapsulate the traffic that it did not originate, and add an RPI in any fashion.

  A packet that is being routed over the RPL Instance associated to a first Non-Storing Mode Track MAY be placed (encapsulated) in a second Track to cover one loose hop of the first Track. On the other hand, a Storing Mode Track must be strict and a packet that it placed in a Storing Mode Track MUST follow that Track till the Track Egress.

  When a Track Egress extracts a packet from a Track (decapsulates the packet), the Destination of the inner packet MUST be either this node or a direct neighbor, or a Target of another Segment of the same Track for which this node is ingress, otherwise the packet MUST be dropped.

All properties of a Track operations are inherited form the main RPL 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.

7.5. Non-Storing Mode Projected Route

As illustrated in Figure 9, a P-DAO that carries an SR-VIO enables the Root to install a source-routed path towards a Track Egress in any particular router.

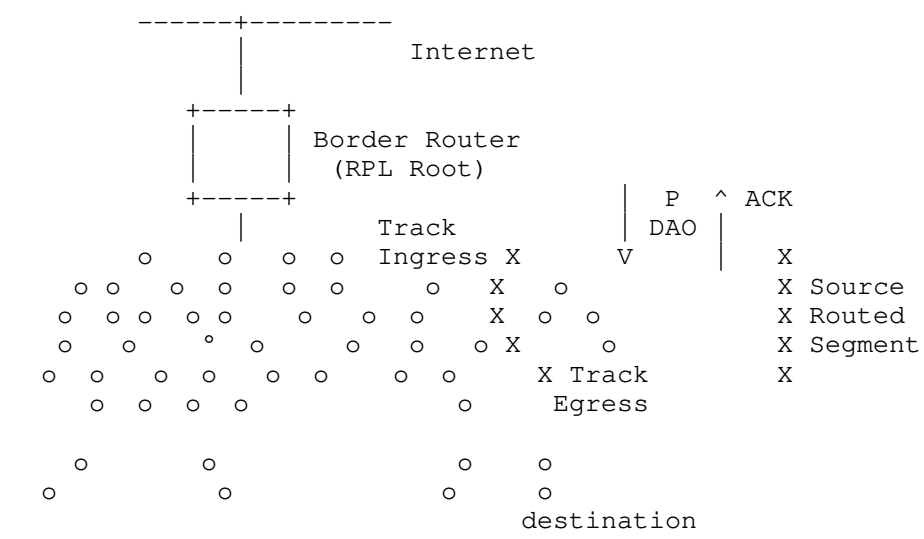

Figure 9: Projecting a Non-Storing Route

A route indicated by an SR-VIO may be loose, meaning that the node that owns the next listed Via Address is not necessarily a neighbor. Without proper loop avoidance mechanisms, the interaction of loose source routing and other mechanisms may effectively cause loops.

When forwarding a packet to a destination for which the router determines that routing happens via the Track Egress, the router inserts the source routing header in the packet with the destination set to the Track Egress.

In order to signal the Segment, the router encapsulates the packet with an IP-in-IP header and a Routing Header as follows:

* In the uncompressed form the source of the packet is this router, the destination is the first Via Address in the SR-VIO, and the RH is a Source Routing Header (SRH) [RFC6554] that contains the list of the remaining Via Addresses terminating by the Track Egress.

* The preferred alternate in a network where 6LoWPAN Header Compression [RFC6282] is used is to leverage "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 source routed header is the exact copy of the (chain of) SRH-6LoRH found in the SR-VIO, also terminating by the Track Egress. The RPI-6LoRH is appended next, followed by an IP-in-IP 6LoRH Header that indicates the Ingress Router in the Encapsulator Address field, see as a similar case Figure 20 of [TURN-ON_RFC8138].

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 another Track to the loose next hop for which this node is Ingress; in the latter case, another encapsulation takes place and the process possibly recurses; otherwise the packet is dropped.

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.

7.6. Storing Mode Projected Route

As illustrated in Figure 10, a P-DAO that carries a SF-VIO enables the Root to install a stateful route towards a collection of Targets along a Segment between a Track Ingress and a Track Egress.

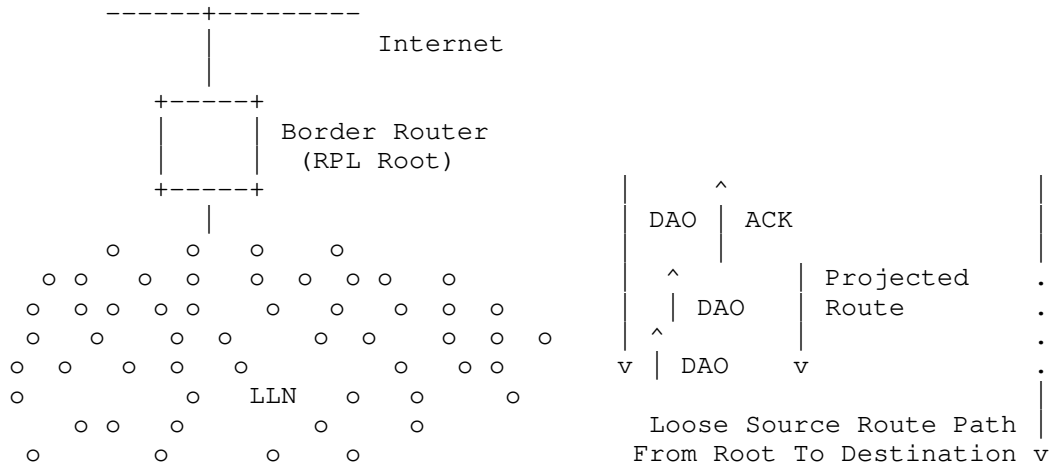

Figure 10: Projecting a route

In order to install the relevant routing state along the Segment , the Root sends a unicast P-DAO message to the Track Egress router of the routing Segment that is being installed. The P-DAO message contains a SF-VIO with the direct sequence of Via Addresses. The SF-VIO follows one or more RTOs indicating the Targets to which the Track leads. The SF-VIO contains a Segment 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 last Via Address in the SF-VIO. This is how the egress recognizes its role. In a similar fashion, the ingress node recognizes its role as it matches first Via Address in the SF-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. If one of the Targets is not known, 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 Address the precedes the one that contain the address of the propagating node, which is used as source of the message.

Upon receiving a propagated DAO, all except the Egress Router MUST install a route towards the DAO Target(s) via their successor in the SF-VIO. The router MAY install additional routes towards the VIA Addresses that are the SF-VIO after the next one, if any, but in case of a conflict or a lack of resource, the route(s) to the Target(s) have precedence.

If a router cannot reach its predecessor in the SF-VIO, 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).

The process continues till the P-DAO is propagated to ingress router of the Segment, which answers with a DAO-ACK to the Root.

A Segment 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 an SMPR, 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.

8. 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?

9. IANA Considerations

9.1. New Elective 6LoWPAN Routing Header Type

This document updates the IANA registry titled "Elective 6LoWPAN Routing Header Type" that was created for [RFC8138] and assigns the following value:

Value	Description	Reference
7	P-RPI-6LoRH	This document

Table 1: New Elective 6LoWPAN Routing Header Type

9.2. New Critical 6LoWPAN Routing Header Type

This document updates the IANA registry titled "Critical 6LoWPAN Routing Header Type" that was created for [RFC8138] and assigns the following value:

Value	Description	Reference
7	P-RPI-6LoRH	This document

Table 2: New Critical 6LoWPAN Routing Header Type

9.3. New Subregistry For The RPL Option Flags

IANA is required to create a subregistry for the 8-bit RPL Option Flags field, as detailed in Figure 2, under the "Routing Protocol for Low Power and Lossy Networks (RPL)" registry. The bits are indexed from 0 (leftmost) to 7. Each bit is tracked with the following qualities:

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

* Indication When Set

* Reference

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

Bit number	Indication When Set	Reference
0	Down 'O'	[RFC6553]
1	Rank-Error (R)	[RFC6553]
2	Forwarding-Error (F)	[RFC6553]
3	Projected-Route (P)	This document

Table 3: Initial PDR Flags

9.4. New RPL Control Codes

This document extends the IANA Subregistry created by RFC 6550 for RPL Control Codes as indicated in Table 4:

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

Table 4: New RPL Control Codes

9.5. New RPL Control Message Options

This document extends the IANA Subregistry created by RFC 6550 for RPL Control Message Options as indicated in Table 5:

Value	Meaning	Reference
0x0B	Stateful Via Information option (SF-VIO)	This document
0x0C	Source-Routed Via Information option (SR-VIO)	This document
0x0D	Sibling Information option	This document

Table 5: RPL Control Message Options

9.6. SubRegistry for the Projected DAO Request 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 6:

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

Table 6: Initial PDR Flags

9.7. 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.

9.8. Subregistry for the PDR-ACK Acceptance Status Values

IANA is requested to create a 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 7:

Value	Meaning	Reference
0	Unqualified acceptance	This document

Table 7: Acceptance values of the PDR-ACK Status

9.9. Subregistry for the PDR-ACK Rejection Status Values

IANA is requested to create a 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 8:

Value	Meaning	Reference
0	Unqualified rejection	This document

Table 8: Rejection values of the PDR-ACK Status

9.10. SubRegistry for the Via Information Options Flags

IANA is requested to create a Subregistry for the 5-bit Via Information Options (Via Option) 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 Via Information Options (Via Option) Flags.

9.11. SubRegistry for the Sibling Information Option 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 9:

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

Table 9: Initial SIO Flags

9.12. 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.

10. Acknowledgments

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

11. 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>.

[RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <https://www.rfc-editor.org/info/rfc6282>.

[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>.

[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>.

[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>.

12. 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, 27 August 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-05, 15 November 2020, <https://tools.ietf.org/html/draft-pthubert-raw-architecture-05>.

[TURN-ON_RFC8138] Thubert, P. and L. Zhao, "A RPL DODAG Configuration Option for the 6LoWPAN Routing Header", Work in Progress, Internet-Draft, draft-ietf-roll-turnon-rfc8138-17, 30 September 2020, <https://tools.ietf.org/html/draft-ietf-roll-turnon-rfc8138-17>.

[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>.

[USEofRPLinfo] Robles, I., Richardson, M., and P. Thubert, "Using RPI Option Type, Routing Header for Source Routes and IPv6-in-IPv6 encapsulation in the RPL Data Plane", Work in Progress, Internet-Draft, draft-ietf-roll-useofrplinfo-42, 12 November 2020, <https://tools.ietf.org/html/draft-ietf-roll-useofrplinfo-42>.

[PCE] IETF, "Path Computation Element", <https://datatracker.ietf.org/doc/charter-ietf-pce/>.

Appendix A. Applications

A.1. Loose Source Routing

A RPL implementation operating in a very constrained LLN typically uses the Non-Storing Mode of Operation as represented in Figure 11. 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.

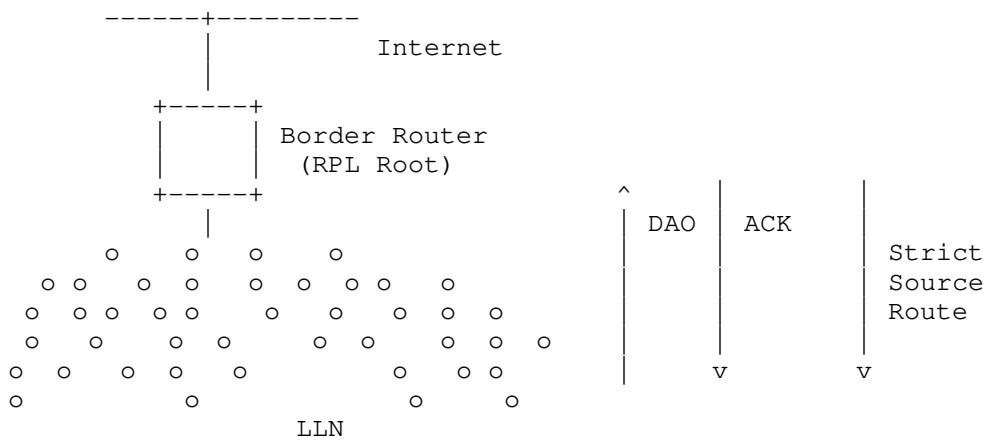

Figure 11: RPL Non-Storing Mode of operation

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 a 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

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 12:

* 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.

* 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 an RH 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.

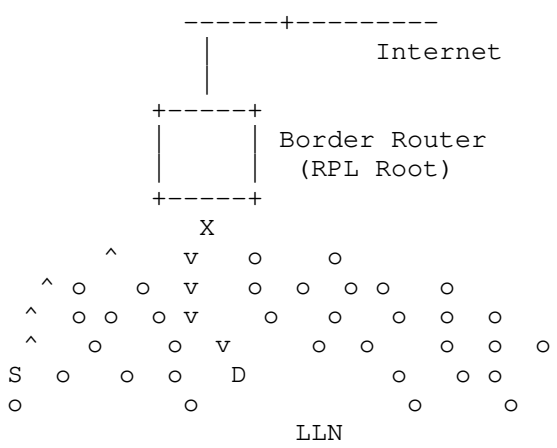

Figure 12: Routing Stretch between S and D via common parent X

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, and that it is even more critical in Non-Storing Mode than it is in Storing Mode, because the routing stretch is wider. 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.

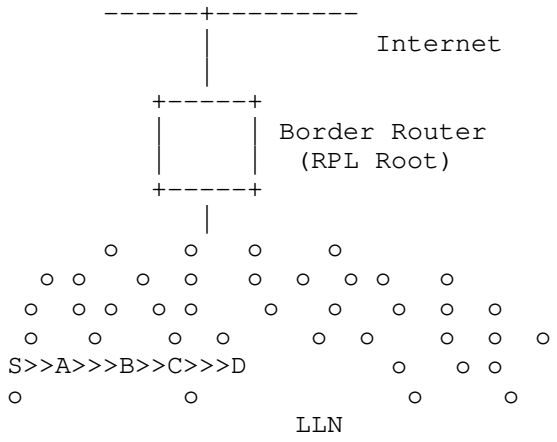

Figure 13: Projected Transversal Route

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 oculd 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.

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
Email: pthubert@cisco.com

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

ROLL
Internet-Draft
Updates: 8138 (if approved)
Intended status: Standards Track
Expires: 3 April 2021

P. Thubert, Ed.
L. Zhao
Cisco Systems

30 September 2020

A RPL DODAG Configuration Option for the 6LoWPAN Routing Header
draft-ietf-roll-turnon-rfc8138-17

Abstract

This document updates RFC 8138 by defining a bit in the RPL DODAG Configuration Option to indicate whether compression is used within the RPL Instance, and specify the behavior of RFC 8138-capable nodes when the bit is set and unset.

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 3 April 2021.

Table of Contents

1. Introduction

The design of Low Power and Lossy Networks (LLNs) is generally focused on saving energy, which is the most constrained resource of all. The routing optimizations in the "Routing Protocol for Low Power and Lossy Networks" [RFC6550] (RPL) such as routing along a Destination-Oriented Directed Acyclic Graph (DODAG) to a Root Node and the associated routing header compression and forwarding technique specified in [RFC8138] derive from that primary concern.

Enabling [RFC8138] on a running network requires a Flag Day where the network is upgraded and rebooted. Otherwise, if acting as a Leaf, a node that does not support the compression would fail to communicate; if acting as a router it would drop the compressed packets and black-hole a portion of the network. This specification enables a hot upgrade where a live network is migrated. During the migration, the compression remains inactive, until all nodes are upgraded.

This document complements [RFC8138] and signals whether it should be used within a RPL DODAG with a new flag in the RPL DODAG Configuration Option. The setting of this new flag is controlled by the Root and propagates as is in the whole network as part of the normal RPL signaling.

The flag is cleared to maintain the compression inactive during the migration phase. When the migration is complete (e.g., as known by network management and/or inventory), the flag is set and the compression is globally activated in the whole DODAG.

2. Terminology

2.1. References

The terminology used in this document is consistent with and incorporates that described in "Terms Used in Routing for Low-Power and Lossy Networks (LLNs)" [RFC7102]. Other terms in use in LLNs are found in "Terminology for Constrained-Node Networks" [RFC7228].

"RPL", the "RPL Packet Information" (RPI), and "RPL Instance" (indexed by a RPLInstanceID) are defined in "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks" [RFC6550]. The RPI is the abstract information that RPL defines to be placed in data packets, e.g., as the RPL Option [RFC6553] within the IPv6 Hop-By-Hop Header. By extension the term "RPI" is often used to refer to the RPL Option itself. The DODAG Information Solicitation (DIS), Destination Advertisement Object (DAO) and DODAG Information Object (DIO) messages are also specified in [RFC6550].

This document uses the terms RPL-Unaware Leaf (RUL) and RPL-Aware Leaf (RAL) consistently with "Using RPI Option Type, Routing Header for Source Routes and IPv6-in-IPv6 encapsulation in the RPL Data Plane" [USEofRPLinfo]. The term RPL-Aware Node (RAN) refers to a node that is either a RAL or a RPL Router. A RAN manages the reachability of its addresses and prefixes by injecting them in RPL by itself. In contrast, a RUL leverages "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery" [RFC8505] to obtain reachability services from its parent router(s) as specified in "Routing for RPL Leaves" [UNAWARE-LEAVES].

2.2. Glossary

This document often uses the following acronyms:

6LoWPAN: IPv6 over Low-Power Wireless Personal Area Network
6LoRH: 6LoWPAN Routing Header
DIO: DODAG Information Object (a RPL message)
DODAG: Destination-Oriented Directed Acyclic Graph
LLN: Low-Power and Lossy Network
RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks
SubDAG: A DODAG rooted at a node which is a child of that node and a subset of a larger DAG
MOP: RPL Mode of Operation
RPI: RPL Packet Information
RAL: RPL-Aware Leaf
RAN: RPL-Aware Node
RUL: RPL-Unaware Leaf
SRH: Source Routing Header

2.3. 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.

3. Extending RFC 6550

The DODAG Configuration Option is defined in Section 6.7.6 of [RFC6550]. Its purpose is extended to distribute configuration information affecting the construction and maintenance of the DODAG, as well as operational parameters for RPL on the DODAG, through the DODAG. As shown in Figure 1, the Option was originally designed with 4 bit positions reserved for future use as Flags.

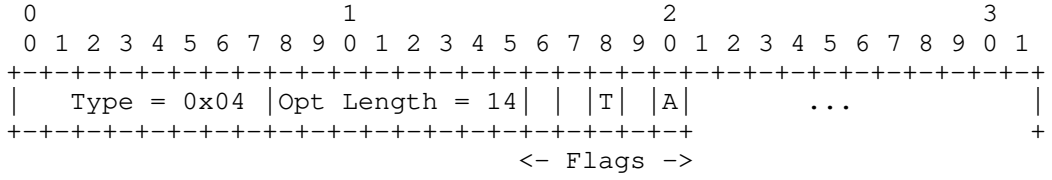

Figure 1: DODAG Configuration Option (Partial View)

This specification defines a new flag "Enable RFC8138 Compression" (T). The "T" flag is set to turn-on the use of [RFC8138] within the DODAG. The "T" flag is encoded in position 2 of the reserved Flags in the DODAG Configuration Option (counting from bit 0 as the most significant bit) and set to 0 in legacy implementations as specified respectively in Sections 20.14 and 6.7.6 of [RFC6550].

Section 4.3 of [USEofRPLinfo] updates [RFC6550] to indicate that the definition of the Flags applies to Mode of Operation (MOP) values zero (0) to six (6) only. For a MOP value of 7, [RFC8138] MUST be used on Links where 6LoWPAN Header Compression [RFC6282] applies and MUST NOT be used otherwise.

The RPL DODAG Configuration Option is typically placed in a DODAG Information Object (DIO) message. The DIO message propagates down the DODAG to form and then maintain its structure. The DODAG Configuration Option is copied unmodified from parents to children. [RFC6550] states that "Nodes other than the DODAG Root MUST NOT modify this information when propagating the DODAG Configuration option". Therefore, a legacy parent propagates the "T" flag as set by the Root, and when the "T" flag is set, it is transparently flooded to all the nodes in the DODAG.

4. Updating RFC 8138

A node SHOULD generate packets in the compressed form using [RFC8138] if and only if the "T" flag is set. This behavior can be overridden by configuration or network management. Overriding may be needed e.g., to turn on the compression in a network where all nodes support [RFC8138] but the Root does not support this specification and cannot set the "T" flag, or to disable it locally in case of a problem.

The decision to use [RFC8138] is made by the originator of the packet depending on its capabilities and its knowledge of the state of the "T" flag. A router encapsulating a packet is the originator of the resulting packet and is responsible for compressing the outer headers with [RFC8138], but it MUST leave the encapsulated packet as is.

An external target [USEofRPLinfo] is not expected to support [RFC8138]. In most cases, packets to and from an external target are tunneled back and forth between the border router (referred to as 6LR) that serves the external target and the Root, regardless of the MOP used in the RPL DODAG. The inner packet is typically not compressed with [RFC8138], so for outgoing packets, the border router just needs to decapsulate the (compressed) outer header and forward the (uncompressed) inner packet towards the external target.

A router MUST uncompress a packet that is to be forwarded to an external target. Otherwise, the router MUST forward the packet in the form that the source used, either compressed or uncompressed.

A RUL [UNAWARE-LEAVES] is both a leaf and an external target. A RUL does not participate in RPL and depends on the parent router to obtain connectivity. In the case of a RUL, forwarding towards an external target actually means delivering the packet.

5. Transition Scenarios

A node that supports [RFC8138] but not this specification can only be used in a homogeneous network. Enabling the [RFC8138] compression without a turn-on signaling method requires a "flag day"; by which time all nodes must be upgraded, and at which point the network can be rebooted with the [RFC8138] compression turned on.

The intent for this specification is to perform a migration once and for all without the need for a flag day. In particular it is not the intention to undo the setting of the "T" flag. Though it is possible to roll back (see Section 5.3), the roll back operation SHOULD be complete before the network operator adds nodes that do not support [RFC8138].

5.1. Coexistence

A node that supports this specification can operate in a network with the [RFC8138] compression turned on or off with the "T" flag set accordingly and in a network in transition from off to on or on to off (see Section 5.2).

A node that does not support [RFC8138] can interoperate with nodes that do in a network with [RFC8138] compression turned off. If the compression is turned on, all the RPL-Aware Nodes are expected to be able to handle compressed packets in the compressed form. A node that cannot do so may remain connected to the network as a RUL as described in [UNAWARE-LEAVES].

5.2. Inconsistent State While Migrating

When the "T" flag is turned on by the Root, the information slowly percolates through the DODAG as the DIO gets propagated. Some nodes will see the flag and start sourcing packets in the compressed form while other nodes in the same RPL DODAG are still not aware of it. In non-storing mode, the Root will start using [RFC8138] with a Source Routing Header 6LoRH (SRH-6LoRH) that routes all the way to the parent router or to the leaf.

To ensure that a packet is forwarded across the RPL DODAG in the form in which it was generated, it is required that all the RPL nodes support [RFC8138] at the time of the switch.

Setting the "T" flag is ultimately the responsibility of the Network Administrator. The expectation is that the network management or upgrading tools in place enable the Network Administrator to know when all the nodes that may join a DODAG were migrated. In the case of a RPL instance with multiple Roots, all nodes that participate to the RPL Instance may potentially join any DODAG. The network MUST be operated with the "T" flag unset until all nodes in the RPL Instance are upgraded to support this specification.

5.3. Rolling Back

When turning [RFC8138] compression off in the network, the Network Administrator MUST wait until all nodes have converged to the "T" flag unset before allowing nodes that do not support the compression in the network. To that effect, whether the compression is active in a node SHOULD be exposed the node's management interface.

Nodes that do not support [RFC8138] SHOULD NOT be deployed in a network where the compression is turned on. If that is done, the node can only operate as a RUL.

6. IANA Considerations

This specification updates the Registry that was created for [RFC6550] as the registry for "DODAG Configuration Option Flags" and updated as the registry for "DODAG Configuration Option Flags for MOP 0..6" by [USEofRPLinfo], by allocating one new Flag as follows:

Bit Number	Capability Description	Reference
2 (suggested)	Turn on RFC8138 Compression (T)	THIS RFC

Table 1: New DODAG Configuration Option Flag

7. Security Considerations

It is worth noting that in RPL [RFC6550], every node in the LLN that is RPL-aware and has access to the RPL domain can inject any RPL-based attack in the network, more in [RFC7416]. This document applies typically to an existing deployment and does not change its security requirements and operations. It is assumed that the security mechanisms as defined for RPL are followed.

Setting the "T" flag before all routers are upgraded may cause a loss of packets. The new bit is protected as the rest of the configuration so this is just one of the many attacks that can happen if an attacker manages to inject a corrupted configuration.

Setting and unsetting the "T" flag may create inconsistencies in the network but as long as all nodes are upgraded to [RFC8138] support they will be able to forward both forms. The source is responsible for selecting whether the packet is compressed or not, and all routers must use the format that the source selected. So the result of an inconsistency is merely that both forms will be present in the network, at an additional cost of bandwidth for packets in the uncompressed form.

An attacker may unset the "T" flag to force additional energy consumption of child or descendant nodes in its subDAG. Conversely it may set the "T" flag, so that nodes located downstream would compress when that it is not desired, potentially resulting in the loss of packets. In a tree structure, the attacker would be in position to drop the packets from and to the attacked nodes. So the attacks above would be more complex and more visible than simply dropping selected packets. The downstream node may have other parents and see both settings, which could raise attention.

8. Acknowledgments

The authors wish to thank Murray Kucherawy, Meral Shirazipour, Barry Leiba, Tirumaleswar Reddy, Nagendra Kumar Nainar, Stewart Bryant, Carles Gomez, Eric Vyncke, Roman Danyliw, and especially Benjamin Kaduk, Alvaro Retana, Dominique Barthel and Rahul Jadhav for their in-depth reviews and constructive suggestions.

Also many thanks to Michael Richardson for being always helpful and responsive when need comes.

9. 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>.

[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>.

[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>.

[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>.

[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>.

[RFC8505] Thubert, P., Ed., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505, DOI 10.17487/RFC8505, November 2018, <https://www.rfc-editor.org/info/rfc8505>.

[UNAWARE-LEAVES]
Thubert, P. and M. Richardson, "Routing for RPL Leaves", Work in Progress, Internet-Draft, draft-ietf-roll-unaware-leaves-18, 12 June 2020, <https://tools.ietf.org/html/draft-ietf-roll-unaware-leaves-18>.

10. Informative References

[RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <https://www.rfc-editor.org/info/rfc6282>.

[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>.

[RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <https://www.rfc-editor.org/info/rfc7228>.

[RFC7416] Tsao, T., Alexander, R., Dohler, M., Daza, V., Lozano, A., and M. Richardson, Ed., "A Security Threat Analysis for the Routing Protocol for Low-Power and Lossy Networks (RPLs)", RFC 7416, DOI 10.17487/RFC7416, January 2015, <https://www.rfc-editor.org/info/rfc7416>.

[USEofRPLinfo]
Robles, I., Richardson, M., and P. Thubert, "Using RPI Option Type, Routing Header for Source Routes and IPv6-in-IPv6 encapsulation in the RPL Data Plane", Work in Progress, Internet-Draft, draft-ietf-roll-useofrplinfo-40, 25 June 2020, <https://tools.ietf.org/html/draft-ietf-roll-useofrplinfo-40>.

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

Email: pthubert@cisco.com

Li Zhao
Cisco Systems, Inc
Xinsi Building
No. 926 Yi Shan Rd
SHANGHAI
200233
China

Email: liz3@cisco.com

ROLL
Internet-Draft
Updates: 6550, 6775, 8505 (if approved)
Intended status: Standards Track
Expires: 14 May 2021

P. Thubert, Ed.
Cisco Systems
M. Richardson
Sandelman

10 November 2020

Routing for RPL Leaves
draft-ietf-roll-unaware-leaves-23

Abstract

This specification updates RFC6550, RFC6775, and RFC8505, to provide routing services to RPL Unaware Leaves that implement 6LoWPAN ND and the extensions therein.

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 14 May 2021.

Copyright Notice

Table of Contents

1. Introduction

The design of Low Power and Lossy Networks (LLNs) is generally focused on saving energy, which is the most constrained resource of all. Other design constraints, such as a limited memory capacity, duty cycling of the LLN devices and low-power lossy transmissions, derive from that primary concern.

The IETF produced the "Routing Protocol for Low Power and Lossy Networks" [RFC6550] (RPL) to provide IPv6 [RFC8200] routing services within such constraints. RPL belongs to the class of Distance-Vector protocols, which, compared to link-state protocols, limit the amount of topological knowledge that needs to be installed and maintained in each node, and does not require convergence to avoid micro-loops.

To save signaling and routing state in constrained networks, RPL allows a path stretch (see [RFC6687]), whereby routing is only performed along a Destination-Oriented Directed Acyclic Graph (DODAG) that is optimized to reach a Root node, as opposed to along the shortest path between 2 peers, whatever that would mean in a given LLN. This trades the quality of peer-to-peer (P2P) paths for a vastly reduced amount of control traffic and routing state that would be required to operate an any-to-any shortest path protocol. Additionally, broken routes may be fixed lazily and on-demand, based on dataplane inconsistency discovery, which avoids wasting energy in the proactive repair of unused paths.

For many of the nodes, though not all, the DODAG provides multiple forwarding solutions towards the Root of the topology via so-called parents. RPL is designed to adapt to fuzzy connectivity, whereby the physical topology cannot be expected to reach a stable state, with a lazy control that creates the routes proactively, but may only fix them reactively, upon actual traffic. The result is that RPL provides reachability for most of the LLN nodes, most of the time, but may not converge in the classical sense.

RPL can be deployed in conjunction with IPv6 Neighbor Discovery (ND) [RFC4861] [RFC4862] and 6LoWPAN ND [RFC6775] [RFC8505] to maintain reachability within a Non-Broadcast Multiple-Access (NBMA) Multi-Link subnet.

In that mode, IPv6 addresses are advertised individually as Host routes. Some nodes may act as Routers and participate in the forwarding operations whereas others will only terminate packets, acting as Hosts in the data-plane. In [RFC6550] terms, an IPv6 Host [RFC8504] that is reachable over the RPL network is called a Leaf.

[USEofRPLinfo] introduces the terms RPL-Aware-Leaf (RAL) and RPL-Unaware Leaf (RUL). A RAL is a Leaf that injects Host routes in RPL to manage the reachability of its IPv6 addresses. Conversely, a RUL does not participate to RPL and cannot inject routes. Section 5 details a Host-to-Router interface that the RUL needs to implement to advertise its IPv6 addresses to a Router that supports this specification. This document specifies how the Router injects those addresses as Host routes in the RPL network on behalf of the RUL.

This specification leverages the Address Registration mechanism defined in 6LoWPAN ND to enable a 6LoWPAN Node (6LN) acting as a RUL to interface with a 6LoWPAN Router (6LR) that is also an RPL-Aware router, and request that this router inject a Host route for the Registered Address in RPL on its behalf. A RUL may be unable to participate because it is very energy-constrained, code-space constrained, or because it would be unsafe to let it inject routes in RPL. Using 6LowPAN ND as the interface for the RUL limits the surface of the possible attacks and optionally protects the address ownership.

The RPL Non-Storing Mode mechanism is used to extend the routing state with connectivity to the RULs even when the DODAG is operated in Storing Mode. The unicast packet forwarding operation by the 6LR serving a RUL is described in section 4.1 of [USEofRPLinfo].

Examples of possible RULs include lightly powered sensors such as window smash sensor (alarm system), and kinetically powered light switches. Other applications of this specification may include a smart grid network that controls appliances - such as washing machines or the heating system - in the home. Appliances may not participate to the RPL protocol operated in the Smartgrid network but can still interact with the Smartgrid for control and/or metering.

This document is organized as follows:

* Section 3 and Section 4 present salient aspects of RPL and 6LoWPAN ND, respectively, that are leveraged in this specification to provide connectivity to a RUL across a RPL network.

* Section 5 lists the expectations that a RUL needs to match in order to be served by a RPL router that complies with this specification.

* Section 6 presents the changes made to [RFC6550]; a new behavior is introduced whereby the 6LR advertises the 6LN's addresses in a RPL DAO message based on the ND registration by the 6LN, and the RPL root performs the EDAR/EDAC exchange with the 6LBR on behalf of the 6LR; modifications are introduced to some RPL options and to the RPL Status to facilitate the integration of the protocols.

* Section 7 presents the changes made to [EFFICIENT-NPDAO]; the use of the DCO message is extended to the Non-Storing MOP to report asynchronous issues from the Root to the 6LR.

* Section 8 presents the changes made to [RFC6775] and [RFC8505]; The range of the ND status codes is reduced down to 64 values, and the remaining bits in the original status field are now reserved.

* Section 9 and Section 10 present the operation of this specification for unicast and multicast flows, respectively, and Section 11 presents associated security considerations.

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. Glossary

This document often uses the following acronyms:

AR: Address Resolution (aka Address Lookup)
ARQ: Automatic Repeat reQuest
6CIO: 6LoWPAN Capability Indication Option
6LN: 6LoWPAN Node (a Low Power Host or Router)
6LR: 6LoWPAN Router
(E)ARO: (Extended) Address Registration Option
(E)DAR: (Extended) Duplicate Address Request
(E)DAC: (Extended) Duplicate Address Confirmation
DAD: Duplicate Address Detection
DAO: Destination Advertisement Object (a RPL message)
DCO: Destination Cleanup Object (a RPL message)
DIS: DODAG Information solicitation (a RPL message)
DIO: DODAG Information Object (a RPL message)
DODAG: Destination-Oriented Directed Acyclic Graph
LLN: Low-Power and Lossy Network

NA: Neighbor Advertisement
NCE: Neighbor Cache Entry
ND: Neighbor Discovery
NS: Neighbor solicitation
RA: Router Advertisement
ROVR: Registration Ownership Verifier
RPI: RPL Packet Information
RAL: RPL-Aware Leaf
RAN: RPL-Aware Node (either a RPL Router or a RPL-Aware Leaf)
RUL: RPL-Unaware Leaf
TID: Transaction ID (a sequence counter in the EARO)

2.3. References

The Terminology used in this document is consistent with and incorporates that described in "Terms Used in Routing for Low-Power and Lossy Networks (LLNs)" [RFC7102]. A glossary of classical 6LoWPAN acronyms is given in Section 2.2. Other terms in use in LLNs are found in "Terminology for Constrained-Node Networks" [RFC7228]. This specification uses the terms 6LN and 6LR to refer specifically to nodes that implement the 6LN and 6LR roles in 6LoWPAN ND and does not expect other functionality such as 6LoWPAN Header Compression [RFC6282] from those nodes.

"RPL", the "RPL Packet Information" (RPI), "RPL Instance" (indexed by a RPLInstanceID) are defined in "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks" [RFC6550]. The RPI is the abstract information that RPL defines to be placed in data packets, e.g., as the RPL Option [RFC6553] within the IPv6 Hop-By-Hop Header. By extension, the term "RPI" is often used to refer to the RPL Option itself. The DODAG Information solicitation (DIS), Destination Advertisement Object (DAO) and DODAG Information Object (DIO) messages are also specified in [RFC6550]. The Destination Cleanup Object (DCO) message is defined in [EFFICIENT-NPDAO].

This document uses the terms RPL-Unaware Leaf (RUL) and RPL Aware Leaf (RAL) consistently with [USEofRPLinfo]. The term RPL-Aware Node (RAN) is introduced to refer to a node that is either an RAL or a RPL Router. As opposed to a RUL, a RAN manages the reachability of its addresses and prefixes by injecting them in RPL by itself.

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

Classical IPv6 ND: "Neighbor Discovery for IP version 6" [RFC4861] and "IPv6 Stateless Address Autoconfiguration" [RFC4862],

6LoWPAN: "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing" [RFC6606] and "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals" [RFC4919], and

6LoWPAN ND: Neighbor Discovery Optimization for Low-Power and Lossy Networks [RFC6775], "Registration Extensions for 6LoWPAN Neighbor Discovery" [RFC8505], and "Address Protected Neighbor Discovery for Low-power and Lossy Networks" [AP-ND].

3. RPL External Routes and Dataplane Artifacts

Section 4.1 of [USEofRPLinfo] provides a set of rules detailed below that must be followed for routing packets from and to a RUL.

A 6LR that acts as a border Router for external routes advertises them using Non-Storing Mode DAO messages that are unicast directly to the Root, even if the DODAG is operated in Storing Mode. Non-Storing Mode routes are not visible inside the RPL domain and all packets are routed via the Root. The RPL Root tunnels the packets directly to the 6LR that advertised the external route, which decapsulates and forwards the original (inner) packet.

The RPL Non-Storing MOP signaling and the associated IP-in-IP encapsulated packets appear as normal traffic to the intermediate Routers. The support of external routes only impacts the Root and the 6LR. It can be operated with legacy intermediate Routers and does not add to the amount of state that must be maintained in those Routers. A RUL is an example of a destination that is reachable via an external route that happens to be also a Host route.

The RPL data packets always carry a Hop-by-Hop Header to transport a RPL Packet Information (RPI) [RFC6550]. So unless the RUL originates its packets with an RPI, the 6LR needs to tunnel them to the Root to add the RPI. As a rule of a thumb and except for the very special case above, the packets from and to a RUL are always encapsulated using an IP-in-IP tunnel between the Root and the 6LR that serves the RUL (see sections 7 and 8 of [USEofRPLinfo] for details). If the packet from the RUL has an RPI, the 6LR as a RPL border router SHOULD rewrite the RPI to indicate the selected Instance and set the flags, but it does not need to encapsulate the packet.

In Non-Storing Mode, packets going down carry a Source Routing Header (SRH). The IP-in-IP encapsulation, the RPI and the SRH are collectively called the "RPL artifacts" and can be compressed using [RFC8138]. Appendix A presents an example compressed format for a packet forwarded by the Root to a RUL in a Storing Mode DODAG.

The inner packet that is forwarded to the RUL may carry some RPL artifacts, e.g., an RPI if the original packet was generated with it, and an SRH in a Non-Storing Mode DODAG. [USEofRPLinfo] expects the RUL to support the basic "IPv6 Node Requirements" [RFC8504]. In particular the RUL is expected to ignore the RPL artifacts that are either consumed or not applicable to a Host.

A RUL is not expected to support the compression method defined in [RFC8138]. For that reason, the border router uncompresses the packet before forwarding over an external route to a RUL [USEofRPLinfo].

4. 6LoWPAN Neighbor Discovery

4.1. RFC 6775 Address Registration

The classical "IPv6 Neighbor Discovery (IPv6 ND) Protocol" [RFC4861] [RFC4862] was defined for serial links and transit media such as Ethernet. It is a reactive protocol that relies heavily on multicast operations for Address Discovery (aka Lookup) and Duplicate Address Detection (DAD).

"Neighbor Discovery Optimizations for 6LoWPAN networks" [RFC6775] adapts IPv6 ND for operations over energy-constrained LLNs. The main functions of [RFC6775] are to proactively establish the Neighbor Cache Entry (NCE) in the 6LR and to prevent address duplication. To that effect, [RFC6775] introduces a new unicast Address Registration mechanism that contributes to reducing the use of multicast messages compared to the classical IPv6 ND protocol.

[RFC6775] defines a new Address Registration Option (ARO) that is carried in the unicast Neighbor solicitation (NS) and Neighbor Advertisement (NA) messages between the 6LoWPAN Node (6LN) and the 6LoWPAN Router (6LR). It also defines the Duplicate Address Request (DAR) and Duplicate Address Confirmation (DAC) messages between the 6LR and the 6LoWPAN Border Router (6LBR). In an LLN, the 6LBR is the central repository of all the Registered Addresses in its domain and the source of truth for uniqueness and ownership.

4.2. RFC 8505 Extended Address Registration

"Registration Extensions for 6LoWPAN Neighbor Discovery" [RFC8505] updates the behavior of RFC 6775 to enable a generic Address Registration to services such as routing and ND proxy, and defines the Extended Address Registration Option (EARO) as shown in Figure 1:

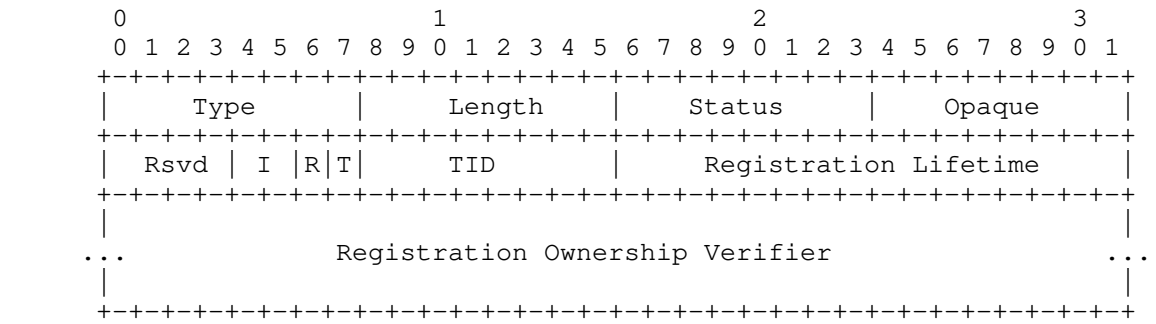

Figure 1: EARO Option Format

4.2.1. R Flag

[RFC8505] introduces the R Flag in the EARO. The Registering Node sets the R Flag to indicate whether the 6LR should ensure reachability for the Registered Address. If the R Flag is not set, then the Registering Node handles the reachability of the Registered Address by other means. In a RPL network, this means that either it is a RAN that injects the route by itself or that it uses another RPL Router for reachability services.

This document specifies how the R Flag is used in the context of RPL. A RPL Leaf that implements the 6LN functionality in [RFC8505] requires reachability services for an IPv6 address if and only if it sets the R Flag in the NS(EARO) used to register the address to a 6LR acting as a RPL border Router. Upon receiving the NS(EARO), the RPL Router generates a DAO message for the Registered Address if and only if the R flag is set.

Section 9.2 specifies additional operations when R flag is set in an EARO that is placed either in an NS or an NA message.

4.2.2. TID, "I" Field and Opaque Fields

When the T Flag is set, the EARO includes a sequence counter called Transaction ID (TID), that is needed to fill the Path Sequence Field in the RPL Transit Option. This is the reason why the support of [RFC8505] by the RUL, as opposed to only [RFC6775] is a prerequisite for this specification (more in Section 5.1). The EARO also transports an Opaque field and an associated "I" field that describes what the Opaque field transports and how to use it.

Section 9.2.1 specifies the use of the "I" field and the Opaque field by a RUL.

4.2.3. ROVR

Section 5.3 of [RFC8505] introduces the Registration Ownership Verifier (ROVR) field of variable length from 64 to 256 bits. The ROVR is a replacement of the EUI-64 in the ARO [RFC6775] that was used to identify uniquely an Address Registration with the Link-Layer address of the owner but provided no protection against spoofing.

"Address Protected Neighbor Discovery for Low-power and Lossy Networks" [AP-ND] leverages the ROVR field as a cryptographic proof of ownership to prevent a rogue third party from registering an address that is already owned. The use of ROVR field enable the 6LR to block traffic that is not sourced at an owned address.

This specification does not address how the protection by [AP-ND] could be extended for use in RPL. On the other hand, it adds the ROVR to the DAO to build the proxied EDAR at the Root (see Section 6.1), which means that nodes that are aware of the Host route are also aware of the ROVR associated to the Target Address.

4.3. RFC 8505 Extended DAR/DAC

[RFC8505] updates the DAR/DAC messages into the Extended DAR/DAC to carry the ROVR field. The EDAR/EDAC exchange takes place between the 6LR and the 6LBR. It is triggered by an NS(EARO) message from a 6LN to create, refresh, and delete the corresponding state in the 6LBR. The exchange is protected by the retry mechanism (ARQ) specified in 8.2.6 of [RFC6775], though in an LLN, a duration longer than the RETRANS_TIMER [RFC4861] of 1 second may be necessary to cover the Turn Around Trip delay between the 6LR and the 6LBR.

RPL [RFC6550] specifies a periodic DAO from the 6LN all the way to the Root that maintains the routing state in the RPL network for the lifetime indicated by the source of the DAO. This means that for each address, there are two keep-alive messages that traverse the whole network, one to the Root and one to the 6LBR.

This specification avoids the periodic EDAR/EDAC exchange across the LLN. The 6LR turns the periodic NS(EARO) from the RUL into a DAO message to the Root on every refresh, but it only generates the EDAR upon the first registration, for the purpose of DAD, which must be verified before the address is injected in RPL. Upon the DAO message, the Root proxies the EDAR exchange to refresh the state at the 6LBR on behalf of the 6LR, as illustrated in Figure 7.

4.3.1. RFC 7400 Capability Indication Option

"6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)" [RFC7400] defines the 6LoWPAN Capability Indication Option (6CIO) that enables a node to expose its capabilities in Router Advertisement (RA) messages.

[RFC8505] defines a number of bits in the 6CIO, in particular:

L: Node is a 6LR.
E: Node is an IPv6 ND Registrar -- i.e., it supports registrations based on EARO.
P: Node is a Routing Registrar, -- i.e., an IPv6 ND Registrar that also provides reachability services for the Registered Address.

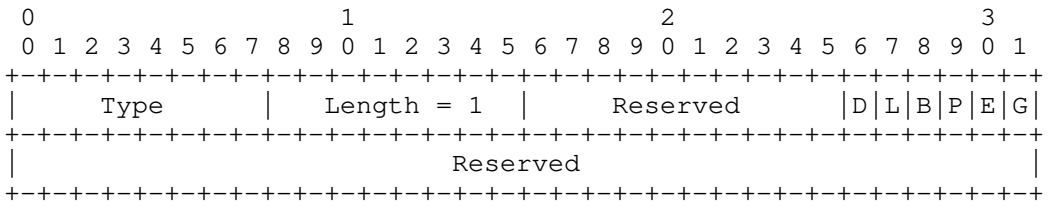

Figure 2: 6CIO flags

A 6LR that can provide reachability services for a RUL in a RPL network as specified in this document MUST include a 6CIO in its RA messages and set the L, P and E flags as prescribed by [RFC8505].

5. Requirements on the RPL-Unware Leaf

This document provides RPL routing for a RUL. This section describes the minimal RPL-independent functionality that the RUL needs to implement to obtain routing services for its addresses.

5.1. Support of 6LoWPAN ND

To obtain routing services from a Router that implements this specification, a RUL needs to implement [RFC8505] and set the "R" and "T" flags in the EARO as discussed in Section 4.2.1 and Section 4.2.3, respectively. Section 9.2.1 specifies new behaviors for the RUL, e.g., when the R Flag set in a NS(EARO) is not echoed in the NA(EARO), which indicates that the route injection failed.

The RUL is expected not to request routing services from a Router that does not originate RA messages with a CIO that has the L, P, and E flags all set as discussed in Section 4.3.1, unless configured to do so. It is suggested that the RUL also implements [AP-ND] to protect the ownership of its addresses.

A RUL that may attach to multiple 6LRs is expected to prefer those that provide routing services. The RUL needs to register to all the 6LRs from which it desires routing services.

Parallel Address Registrations to several 6LRs should be performed in a rapid sequence, using the same EARO for the same Address. Gaps between the Address Registrations will invalidate some of the routes till the Address Registration finally shows on those routes.

[RFC8505] introduces error Status values in the NA(EARO) which can be received synchronously upon an NS(EARO) or asynchronously. The RUL needs to support both cases and refrain from using the address when the Status value indicates a rejection (see Section 6.3).

5.2. Support of IPv6 Encapsulation

Section 2.1 of [USEofRPLinfo] defines the rules for tunneling either to the final destination (e.g., a RUL) or to its attachment Router (designated as 6LR). To terminate the IP-in-IP tunnel, the RUL, as an IPv6 Host, must be able to decapsulate the tunneled packet and either drop the inner packet if it is not the final destination, or pass it to the upper layer for further processing. Unless it is aware by other means that the RUL can handle IP-in-IP properly, which is not mandated by [RFC8504], the Root terminates the IP-in-IP tunnel at the parent 6LR. It is thus not necessary for a RUL to support IP-in-IP decapsulation.

5.3. Support of the HbH Header

A RUL is expected to process an Option Type in a Hop-by-Hop Header as prescribed by section 4.2 of [RFC8200]. An RPI with an Option Type of 0x23 [USEofRPLinfo] is thus skipped when not recognized.

5.4. Support of the Routing Header

A RUL is expected to process an unknown Routing Header Type as prescribed by section 4.4 of [RFC8200]. This implies that the Source Routing Header with a Routing Type of 3 [RFC6554] is ignored when the Segments Left is zero, and the packet is dropped otherwise.

6. Enhancements to RFC 6550

This document specifies a new behavior whereby a 6LR injects DAO messages for unicast addresses (see Section 9) and multicast addresses (see Section 10) on behalf of leaves that are not aware of RPL. The RUL addresses are exposed as external targets [RFC6550]. Conforming to [USEofRPLinfo], an IP-in-IP encapsulation between the 6LR and the RPL Root is used to carry the RPL artifacts and remove them when forwarding outside the RPL domain, e.g., to a RUL.

This document also synchronizes the liveness monitoring at the Root and the 6LBR. The same value of lifetime is used for both, and a single keep-alive message, the RPL DAO, traverses the RPL network. A new behavior is introduced whereby the RPL Root proxies the EDAR message to the 6LBR on behalf of the 6LR (more in Section 8), for any Leaf node that implements the 6LN functionality in [RFC8505].

Section 6.7.7 of [RFC6550] introduces the RPL Target Option, which can be used in RPL Control messages such as the DAO message to signal a destination prefix. This document adds the capabilities to transport the ROVR field (see Section 4.2.3) and the IPv6 Address of the prefix advertiser when the Target is a shorter prefix. Their use is signaled respectively by a new ROVR Size field being non-zero and a new "Advertiser address in Full" 'F' flag set, more in Section 6.1.

This specification defines the new "Root Proxies EDAR/EDAC" (P) flag and encodes it in one of these reserved flags of the RPL DODAG Configuration option, more in Section 6.2.

The RPL Status defined in section 6.5.1 of [RFC6550] for use in the DAO-ACK message is extended to be placed in DCO messages [EFFICIENT-NPDAO] as well. Furthermore, this specification enables to carry the EARO Status defined for 6LoWPAN ND in RPL DAO and DCO messages, embedded in a RPL Status, more in Section 6.3.

Section 12 of [RFC6550] details the RPL support for multicast flows when the RPLInstance is operated in the MOP of 3 ("Storing Mode of Operation with multicast support"). This specification extends the RPL Root operation to proxy-relay the MLDv2 [RFC3810] operation between the RUL and the 6LR, more in Section 10.

6.1. Updated RPL Target Option

This specification updates the RPL Target Option to transport the ROVR that was also defined for 6LoWPAN ND messages. This enables the RPL Root to generate the proxied EDAR message to the 6LBR.

The new 'F' flag is set to indicate that the Target Prefix field contains the IPv6 address of the advertising node, in which case the length of the Target Prefix field is 128 bits regardless of the value of the Prefix Length field. If the 'F' flag is reset, the Target Prefix field MUST be aligned to the next byte boundary after the size (expressed in bits) indicated by the Prefix Length field. Padding bits are reserved and set to 0 per section 6.7.7 of [RFC6550].

With this specification the ROVR is the remainder of the RPL Target Option. The size of the ROVR is indicated in a new ROVR Size field that is encoded to map one-to-one with the Code Suffix in the EDAR message (see table 4 of [RFC8505]). The ROVR Size field is taken from the flags field, which is an update to the RPL Target Option Flags IANA registry.

The updated format is illustrated in Figure 3. It is backward compatible with the Target Option in [RFC6550]. It is recommended that the updated format be used as a replacement in new implementations in all MOPs in preparation for upcoming Route Ownership Validation mechanisms based on the ROVR, unless the device or the network is so constrained that this is not feasible.

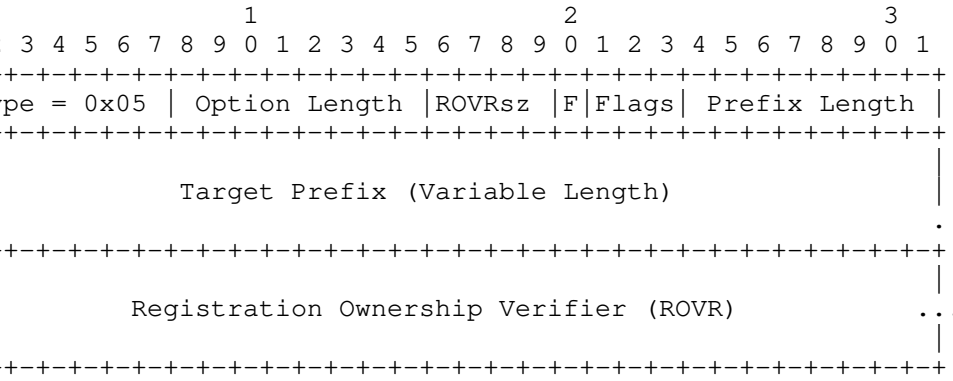

Figure 3: Updated Target Option

New fields:

ROVRsz (ROVR Size): Indicates the Size of the ROVR. It SHOULD be 1, 2, 3, or 4, indicating a ROVR size of 64, 128, 192, or 256 bits, respectively. If a legacy Target Option is used, then the value must remain 0, as specified in [RFC6550]. In case of a value above 4, the size of the ROVR is undetermined and this node cannot validate the ROVR; an implementation SHOULD propagate the whole Target Option upwards as received to enable the verification by an ancestor that would support the upgraded ROVR.

F: 1-bit flag. Set to indicate that Target Prefix field contains the complete (128 bit) IPv6 address of the advertising node.

Registration Ownership Verifier (ROVR): This is the same field as in the EARO, see [RFC8505]

6.2. New Flag in the RPL DODAG Configuration Option

The DODAG Configuration Option is defined in Section 6.7.6 of [RFC6550]. Its purpose is extended to distribute configuration information affecting the construction and maintenance of the DODAG, as well as operational parameters for RPL on the DODAG, through the DODAG. This Option was originally designed with 4 bit positions reserved for future use as Flags.

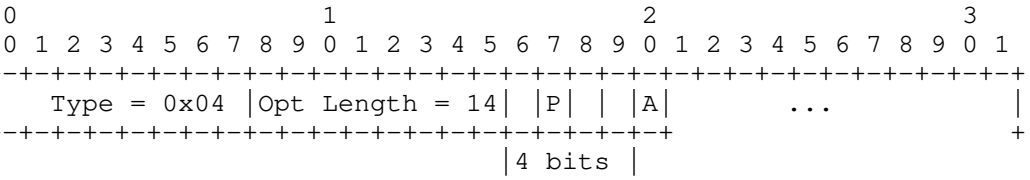

Figure 4: DODAG Configuration Option (Partial View)

This specification defines a new flag "Root Proxies EDAR/EDAC" (P). The 'P' flag is encoded in bit position 1 of the reserved Flags in the DODAG Configuration Option (counting from bit 0 as the most significant bit) and it is set to 0 in legacy implementations as specified respectively in Sections 20.14 and 6.7.6 of [RFC6550].

The 'P' flag is set to indicate that the Root performs the proxy operation, which implies that it supports this specification and the updated RPL Target Option (see Section 6.1).

Section 4.3 of [USEofRPLinfo] updates [RFC6550] to indicate that the definition of the Flags applies to Mode of Operation (MOP) values zero (0) to six (6) only. For a MOP value of 7, the implementation MUST consider that the Root performs the proxy operation.

The RPL DODAG Configuration Option is typically placed in a DODAG Information Object (DIO) message. The DIO message propagates down the DODAG to form and then maintain its structure. The DODAG Configuration Option is copied unmodified from parents to children. [RFC6550] states that "Nodes other than the DODAG Root MUST NOT modify this information when propagating the DODAG Configuration option". Therefore, a legacy parent propagates the 'P' Flag as set by the Root, and when the 'P' Flag is set, it is transparently flooded to all the nodes in the DODAG.

6.3. Updated RPL Status

The RPL Status is defined in section 6.5.1 of [RFC6550] for use in the DAO-ACK message and values are assigned as follows:

Range	Meaning
0	Success/Unqualified acceptance
1-127	Not an outright rejection
128-255	Rejection

Table 1: RPL Status per RFC 6550

The 6LoWPAN ND Status was defined for use in the EARO, see section 4.1 of [RFC8505]. This specification enables to carry the 6LoWPAN ND Status values in RPL DAO and DCO messages, embedded in the RPL Status field.

To achieve this, the range of the ARO/EARO Status values is reduced to 0-63, which updates the IANA registry created for [RFC6775]. This reduction ensures that the values fit within a RPL Status as shown in Figure 5. See Section 12.2, Section 12.5, and Section 12.6 for the respective IANA declarations.

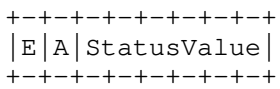

Figure 5: RPL Status Format

This specification updates the RPL Status with subfields as indicated below:

E: 1-bit flag. Set to indicate a rejection. When not set, a Status value of 0 indicates Success/Unqualified acceptance and other values indicate "not an outright rejection" as per RFC 6550.

A: 1-bit flag. Indicates the type of the RPL Status value.

Status Value: 6-bit unsigned integer. If the 'A' flag is set this field transports a Status value defined for IPv6 ND EARO. When the 'A' flag is not set, the Status value is defined for RPL.

When building a DCO or a DAO-ACK message upon an IPv6 ND NA or a EDAC message, the RPL Root MUST copy the 6LoWPAN ND status code unchanged in the RPL Status value and set the 'A' flag. The RPL Root MUST set the 'E' flag for all rejection and unknown status codes. The status codes in the 1-10 range [RFC8505] are all considered rejections.

Reciprocally, upon a DCO or a DAO-ACK message from the RPL Root with a RPL Status that has the 'A' flag set, the 6LR MUST copy the RPL Status value unchanged in the Status field of the EARO when generating an NA to the RUL.

7. Enhancements to draft-ietf-roll-efficient-npdao

[EFFICIENT-NPDAO] defines the DCO message for RPL Storing Mode only, with a link-local scope. All nodes in the RPL network are expected to support the specification since the message is processed hop by hop along the path that is being cleaned up.

This specification extends the use of the DCO message to the Non-Storing MOP, whereby the DCO is sent end-to-end by the Root directly to the RAN that injected the DAO message for the considered target. In that case, intermediate nodes do not need to support [EFFICIENT-NPDAO]; they forward the DCO message as a plain IPv6 packet between the Root and the RAN.

This specification leverages the Non-Storing DCO between the Root and the 6LR that serves as attachment Router for a RUL. A 6LR and a Root that support this specification MUST implement the Non-Storing DCO.

8. Enhancements to RFC 6775 and RFC8505

This document updates [RFC6775] and [RFC8505] to reduce the range of the ND status codes down to 64 values. The two most significant (leftmost) bits if the original ND status field are now reserved, they MUST be set to zero by the sender and ignored by the receiver.

This document also changes the behavior of a 6LR acting as RPL Router and of a 6LN acting as RUL in the 6LoWPAN ND Address Registration as follows:

* If the RPL Root advertises the capability to proxy the EDAR/EDAC exchange to the 6LBR, the 6LR refrains from sending the keep-alive EDAR message. If it is separated from the 6LBR, the Root regenerates the EDAR message to the 6LBR periodically, upon a DAO message that signals the liveliness of the address.

* The use of the R Flag is extended to the NA(EARO) to confirm whether the route was installed.

9. Protocol Operations for Unicast Addresses

The description below assumes that the Root sets the 'P' flag in the DODAG Configuration Option and performs the EDAR proxy operation.

If the 'P' flag is reset, the 6LR MUST generate the periodic EDAR messages and process the returned status as specified in [RFC8505]. If the EDAC indicates success, the rest of the flow takes place as presented but without the proxied EDAR/EDAC exchange.

Section 9.1 provides an overview of the route injection in RPL, whereas Section 9.2 offers more details from the perspective of the different nodes involved in the flow.

9.1. General Flow

This specification eliminates the need to exchange keep-alive Extended Duplicate Address messages, EDAR and EDAC, all the way from a 6LN to the 6LBR across a RPL mesh. Instead, the EDAR/EDAC exchange with the 6LBR is proxied by the RPL Root upon the DAO message that refreshes the RPL routing state. The first EDAR upon a new Registration cannot be proxied, though, as it serves for the purpose of DAD, which must be verified before the address is injected in RPL.

In a RPL network where the function is enabled, refreshing the state in the 6LBR is the responsibility of the Root. Consequently, only addresses that are injected in RPL will be kept alive at the 6LBR by the RPL Root. Since RULs are advertised using Non-Storing Mode, the DAO message flow and the keep alive EDAR/EDAC can be nested within the Address (re)Registration flow. Figure 6 illustrates that, for the first Registration, both the DAD and the keep-alive EDAR/EDAC exchanges happen in the same sequence.

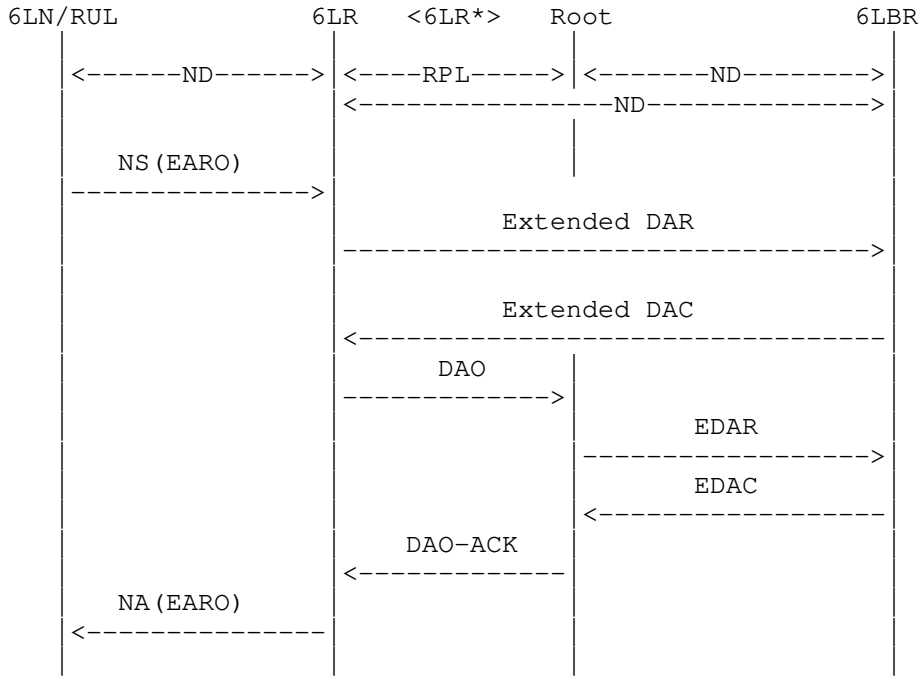

Figure 6: First RUL Registration Flow

This flow requires that the lifetimes and sequence counters in 6LoWPAN ND and RPL are aligned.

To achieve this, the Path Sequence and the Path Lifetime in the DAO message are taken from the Transaction ID and the Address Registration lifetime in the NS(EARO) message from the 6LN.

On the first Address Registration, illustrated in Figure 6 for RPL Non-Storing Mode, the Extended Duplicate Address exchange takes place as prescribed by [RFC8505]. If the exchange fails, the 6LR returns an NA message with a negative status to the 6LN, the NCE is not created, and the address is not injected in RPL. Otherwise, the 6LR creates an NCE and injects the Registered Address in the RPL routing using a DAO/DAO-ACK exchange with the RPL DODAG Root.

An Address Registration refresh is performed by the 6LN to maintain the NCE in the 6LR alive before the lifetime expires. Upon the refresh of a registration, the 6LR reinjects the corresponding route in RPL before it expires, as illustrated in Figure 7.

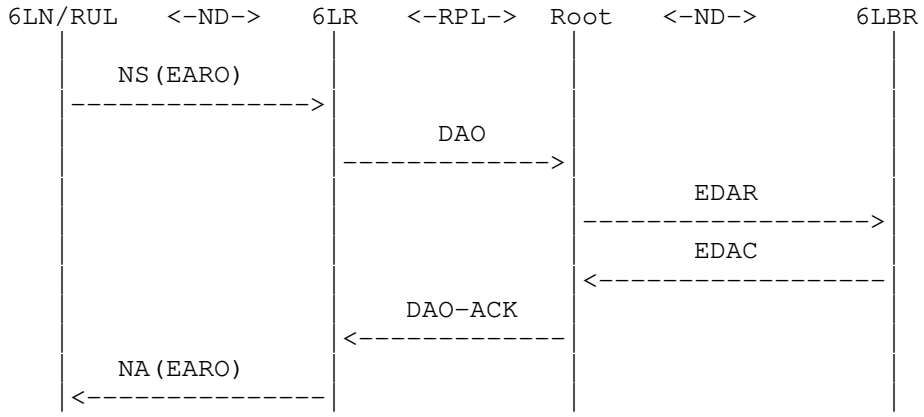

Figure 7: Next RUL Registration Flow

This is what causes the RPL Root to refresh the state in the 6LBR, using an EDAC message. In case of an error in the proxied EDAR flow, the error is returned in the DAO-ACK using a RPL Status with the 'A' flag set that imbeds a 6LoWPAN Status value as discussed in Section 6.3.

The 6LR may receive a requested DAO-ACK after it received an asynchronous DCO, but the negative Status in the DCO supersedes a positive Status in the DAO-ACK regardless of the order in which they are received. Upon the DAO-ACK - or the DCO if one arrives first - the 6LR responds to the RUL with an NA(EARO).

An issue may be detected later, e.g., the address moves to a different DODAG with the 6LBR attached to a different 6LoWPAN Backbone Router (6BBR), see Figure 5 in section 3.3 of [6BBR]. The 6BBR may send a negative ND status, e.g., in an asynchronous NA(EARO) to the 6LBR.

[6BBR] expects that the 6LBR is collocated with the RPL Root, but if not, the 6LBR MUST forward the status code to the originator of the EDAR, either the 6LR or the RPL Root that proxies for it. The ND status code is mapped in a RPL Status value by the RPL Root, and then back by the 6LR.

Figure 8 illustrates this in the case where the 6LBR and the Root are not collocated, and the Root proxies the EDAR messages.

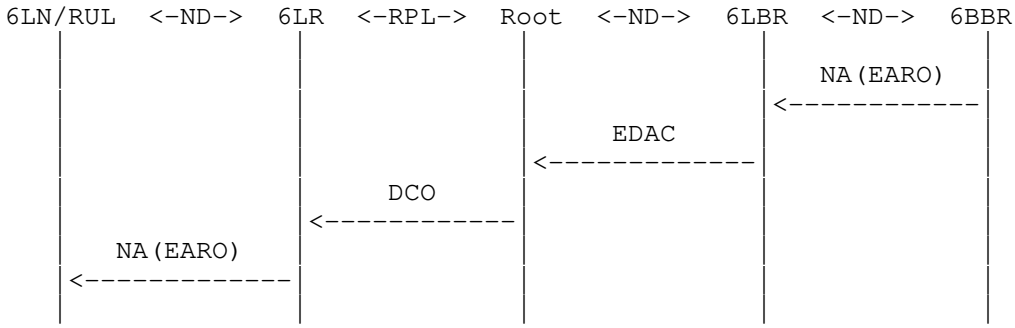

Figure 8: Asynchronous Issue

If the Root does not proxy, then the EDAC with a negative status reaches the 6LR directly. In that case, the 6LR MUST clean up the route using a DAO with a Lifetime of zero, and it MUST propagate the status back to the RUL in a NA(EARO) with the R Flag not set.

The RUL may terminate the registration at any time by using a Registration Lifetime of 0. This specification requires that the RPL Target Option transports the ROVR. This way, the same flow as the heartbeat flow is sufficient to inform the 6LBR using the Root as proxy, as illustrated in Figure 7.

Any combination of the logical functions of 6LR, Root, and 6LBR might be collapsed in a single node.

9.2. Detailed Operation

9.2.1. Perspective of the 6LN Acting as RUL

This specification does not alter the operation of a 6LoWPAN ND-compliant 6LN/RUL, which is expected to operate as follows:

1. The 6LN obtains an IPv6 global address, either using Stateless Address Autoconfiguration (SLAAC) [RFC4862] based on a Prefix Information Option (PIO) [RFC4861] found in an RA message, or some other means, such as DHCPv6 [RFC8415].

2. Once it has formed an address, the 6LN registers its address and refreshes its registration periodically, early enough within the Lifetime of the previous Address Registration, as prescribed by [RFC6775], to refresh the NCE before the lifetime indicated in the EARO expires. It MUST set the T Flag. The TID is incremented each time and wraps in a lollipop fashion (see section 5.2.1 of [RFC8505], which is fully compatible with section 7.2 of [RFC6550]).

3. As stated in section 5.2 of [RFC8505], the 6LN can register to more than one 6LR at the same time. In that case, it uses the same EARO for all of the parallel Address Registrations, with the exception of the Registration Lifetime field and the setting of the R flag that may differ. The 6LN may cancel a subset of its registrations, or transfer a registration from one or more old 6LR(s) to one or more new 6LR(s). To do so, the 6LN sends a series of NS(EARO) messages, all with the same TID, with a zero Registration Lifetime to the old 6LR(s) and with a non-zero Registration Lifetime to the new 6LR(s). In that process, the 6LN SHOULD send the NS(EARO) with a non-zero Registration Lifetime and ensure that at least one succeeds before it sends an NS(EARO) that terminates another registration. This avoids the churn related to transient route invalidation in the RPL network above the common parent of the involved 6LRs.

4. Following section 5.1 of [RFC8505], a 6LN acting as a RUL sets the R Flag in the EARO of its registration(s) for which it requires routing services. If the R Flag is not echoed in the NA, the RUL SHOULD attempt to use another 6LR. The RUL SHOULD ensure that one registration succeeds before resetting the R Flag. In case of a conflict with the preceding rule on lifetime, the rule on lifetime has precedence.

5. The 6LN may use any of the 6LRs to which it registered as the default gateway. Using a 6LR to which the 6LN is not registered may result in packets dropped at the 6LR by a Source Address Validation function (SAVI) [RFC7039] so it is not recommended.

Even without support for RPL, the RUL may be configured with an opaque value to be provided to the routing protocol. If the RUL has knowledge of the RPL Instance the packet should be injected into, then it SHOULD set the Opaque field in the EARO to the RPLInstanceID, else it MUST leave the Opaque field to zero.

Regardless of the setting of the Opaque field, the 6LN MUST set the "I" field to zero to signal "topological information to be passed to a routing process", as specified in section 5.1 of [RFC8505].

A RUL is not expected to produce RPL artifacts in the data packets, but it may do so. For instance, if the RUL has minimal awareness of the RPL Instance then it can build an RPI. A RUL that places an RPI in a data packet SHOULD indicate the RPLInstanceID of the RPL Instance where the packet should be forwarded. It is up to the 6LR (e.g., by policy) to use the RPLInstanceID information provided by the RUL or rewrite it to the selected RPLInstanceID for forwarding inside the RPL domain. All the flags and the Rank field are set to zero as specified by section 11.2 of [RFC6550].

9.2.2. Perspective of the 6LR Acting as Border Router

As prescribed by [RFC8505], the 6LR generates an EDAR message upon reception of a valid NS(EARO) message for the registration of a new IPv6 address by a 6LN. If the initial EDAR/EDAC exchange succeeds, then the 6LR installs an NCE for the Registration Lifetime. For the registration refreshes, if the RPL Root has indicated that it proxies the keep-alive EDAR/EDAC exchange with the 6LBR (see Section 6), the 6LR MUST refrain from sending the keep-alive EDAR.

If the R Flag is set in the NS(EARO), the 6LR SHOULD inject the Host route in RPL, unless this is barred for other reasons, such as the saturation of the RPL parents. The 6LR MUST use a RPL Non-Storing Mode signaling and the updated Target Option (see Section 6.1). The 6LR MUST request a DAO-ACK by setting the 'K' flag in the DAO message. Success injecting the route to the RUL's address is indicated by the 'E' flag set to 0 in the RPL status of the DAO-ACK message.

The Opaque field in the EARO provides a mean to signal which RPL Instance is to be used for the DAO advertisements and the forwarding of packets sourced at the Registered Address when there is no RPI in the packet.

As described in [RFC8505], if the "I" field is zero, then the Opaque field is expected to carry the RPLInstanceID suggested by the 6LN; otherwise, there is no suggested Instance. If the 6LR participates in the suggested RPL Instance, then the 6LR MUST use that RPL Instance for the Registered Address.

If there is no suggested RPL Instance or else if the 6LR does not participate to the suggested Instance, it is expected that the packets coming from the 6LN "can unambiguously be associated to at least one RPL Instance" [RFC6550] by the 6LR, e.g., using a policy that maps the 6-tuple into an Instance.

The DAO message advertising the Registered Address MUST be constructed as follows:

1. The Registered Address is signaled as the Target Prefix in the updated Target Option in the DAO message; the Prefix Length is set to 128 but the 'F' flag is not set since the advertiser is not the RUL. The ROVR field is copied unchanged from the EARO (see Section 6.1).

2. The 6LR indicates one of its global or unique-local IPv6 unicast addresses as the Parent Address in the RPL Transit Information Option (TIO) associated with the Target Option

3. The 6LR sets the External 'E' flag in the TIO to indicate that it is redistributing an external target into the RPL network

4. the Path Lifetime in the TIO is computed from the Registration Lifetime in the EARO. This operation converts seconds to the Lifetime Units used in the RPL operation. This creates the deployment constraint that the Lifetime Unit is reasonably compatible with the expression of the Registration Lifetime. e.g., a Lifetime Unit of 0x4000 maps the most significant byte of the Registration Lifetime to the Path Lifetime.

   In that operation, the Path Lifetime must be rounded, if needed, to the upper value to ensure that the path has a longer lifetime than the registration.

   Note that if the Registration Lifetime is 0, then the Path Lifetime is also 0 and the DAO message becomes a No-Path DAO, which cleans up the routes down to the RUL's address; this also causes the Root as a proxy to send an EDAR message to the 6LBR with a Lifetime of 0.

5. the Path Sequence in the TIO is set to the TID value found in the EARO option.

Upon receiving or timing out the DAO-ACK after an implementation-specific number of retries, the 6LR MUST send the corresponding NA(EARO) to the RUL. Upon receiving an asynchronous DCO message, if a DAO-ACK is pending then the 6LR MUST wait for the DAO-ACK to send the NA(EARO) and deliver the status found in the DCO, else it MUST send an asynchronous NA(EARO) to the RUL immediately.

The 6LR MUST set the R Flag in the NA(EARO) back if and only if the 'E' flag is reset, indicating that the 6LR injected the Registered Address in the RPL routing successfully and that the EDAR proxy operation succeeded.

If the 'A' flag in the RPL Status is set, the embedded Status value is passed back to the RUL in the EARO Status. If the 'E' flag is also set, the registration failed for 6LoWPAN ND related reasons, and the NCE is removed.

An error injecting the route causes the 'E' flag to be set. If the error is not related to ND, the 'A' flag is not set. In that case, the registration succeeds, but the RPL route is not installed. So the NA(EARO) is returned with a positive status but the R Flag not set, which means that the 6LN obtained a binding but no route.

If the 'A' flag is not set in the RPL Status of the DAO-ACK, then the 6LoWPAN ND operation succeeded, and an EARO Status of 0 (Success) MUST be returned to the 6LN. The EARO Status of 0 MUST also be used if the 6LR did not attempt to inject the route but could create the binding after a successful EDAR/EDAC exchange or refresh it.

If the 'E' flag is set in the RPL Status of the DAO-ACK, then the route was not installed and the R flag MUST NOT be set in the NA(EARO). The R flag MUST NOT be set if the 6LR did not attempt to inject the route.

In a network where Address Protected Neighbor Discovery (AP-ND) is enabled, in case of a DAO-ACK or a DCO indicating transporting an EARO Status value of 5 (Validation Requested), the 6LR MUST challenge the 6LN for ownership of the address, as described in section 6.1 of [AP-ND], before the Registration is complete. This flow, illustrated in Figure 9, ensures that the address is validated before it is injected in the RPL routing.

If the challenge succeeds, then the operations continue as normal. In particular, a DAO message is generated upon the NS(EARO) that proves the ownership of the address. If the challenge failed, the 6LR rejects the registration as prescribed by AP-ND and may take actions to protect itself against DoS attacks by a rogue 6LN, see Section 11.

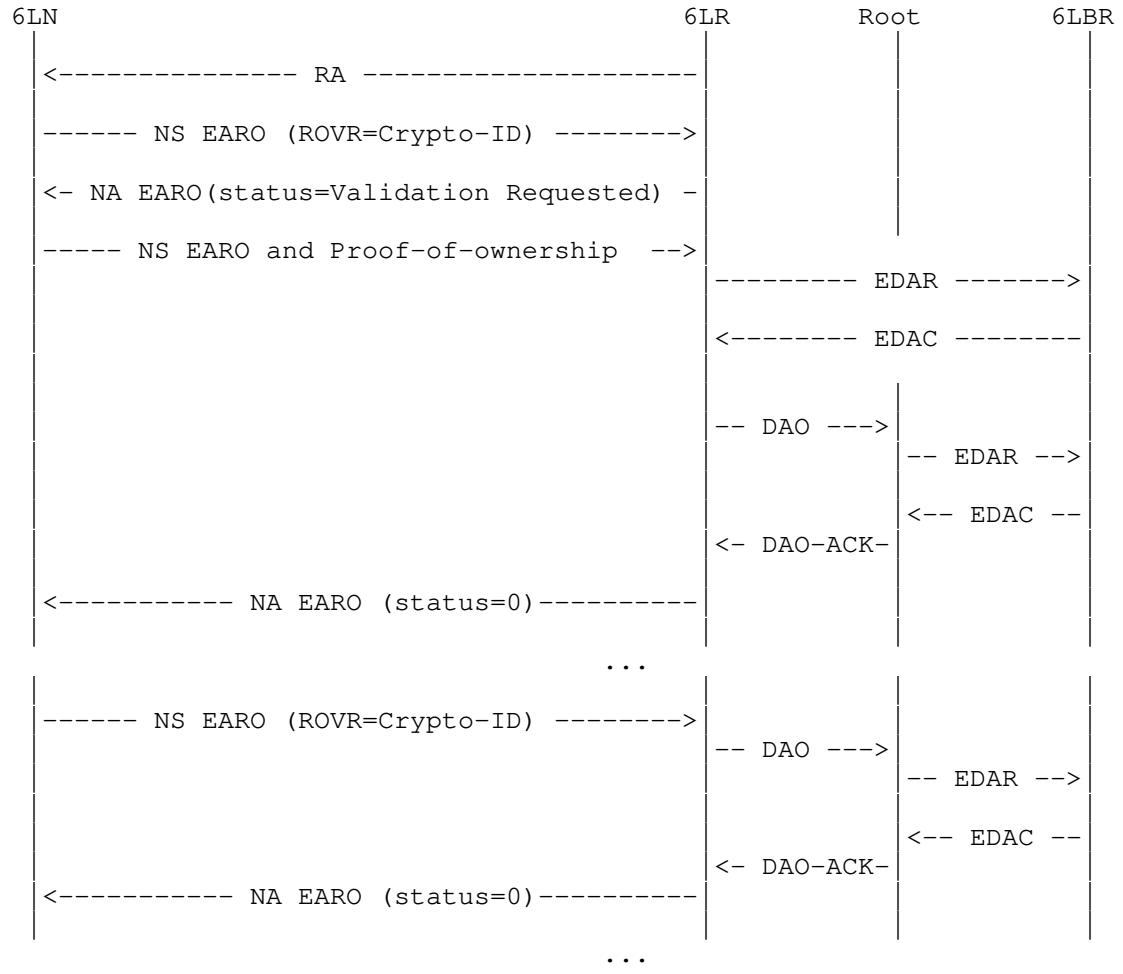

Figure 9: Address Protection

The 6LR may at any time send a unicast asynchronous NA(EARO) with the R Flag reset to signal that it stops providing routing services, and/or with the EARO Status 2 "Neighbor Cache full" to signal that it removes the NCE. It may also send a final RA, unicast or multicast, with a Router Lifetime field of zero, to signal that it stops serving as Router, as specified in section 6.2.5 of [RFC4861]. This may happen upon a DCO or a DAO-ACK message indicating the path is already removed; else the 6LR MUST remove the Host route to the 6LN using a DAO message with a Path Lifetime of zero.

A valid NS(EARO) message with the R Flag not set and a Registration Lifetime that is not zero signals that the 6LN wishes to maintain the binding but does not require the routing services from the 6LR (any more). Upon this message, if, due to previous NS(EARO) with the R Flag set, the 6LR was injecting the Host route to the Registered Address in RPL using DAO messages, then the 6LR MUST invalidate the Host route in RPL using a DAO with a Path Lifetime of zero. It is up to the Registering 6LN to maintain the corresponding route from then on, either keeping it active via a different 6LR or by acting as a RAN and managing its own reachability.

9.2.3. Perspective of the RPL Root

A RPL Root MUST set the 'P' flag in the RPL DODAG Configuration Option of the DIO messages that it generates (see Section 6) to signal that it proxies the EDAR/EDAC exchange and supports the Updated RPL Target option.

Upon reception of a DAO message, for each updated RPL Target Option (see Section 6.1) that creates or updates an existing RPL state, the Root MUST notify the 6LBR by using a proxied EDAR/EDAC exchange. If if the RPL Root and the 6LBR are integrated, an internal API can be used.

The EDAR message MUST be constructed as follows:

1. The Target IPv6 address from the RPL Target Option is placed in the Registered Address field of the EDAR message;

2. the Registration Lifetime is adapted from the Path Lifetime in the TIO by converting the Lifetime Units used in RPL into units of 60 seconds used in the 6LoWPAN ND messages;

3. the TID value is set to the Path Sequence in the TIO and indicated with an ICMP code of 1 in the EDAR message;

4. The ROVR in the RPL Target Option is copied as is in the EDAR and the ICMP Code Suffix is set to the appropriate value as shown in Table 4 of [RFC8505] depending on the size of the ROVR field.

Upon receiving an EDAC message from the 6LBR, if a DAO is pending, then the Root MUST send a DAO-ACK back to the 6LR. Else, if the Status in the EDAC message is not "Success", then it MUST send an asynchronous DCO to the 6LR.

In either case, the EDAC Status is embedded in the RPL Status with the 'A' flag set.

The proxied EDAR/EDAC exchange MUST be protected with a timer of an appropriate duration and a number of retries, that are implementation-dependent, and SHOULD be configurable since the Root and the 6LBR are typically nodes with a higher capacity and manageability than 6LRs. Upon timing out, the Root MUST send an error back to the 6LR as above, either using a DAO-ACK or a DCO, as appropriate, with the 'A' and 'E' flags set in the RPL status, and a RPL Status value of of "6LBR Registry Saturated" [RFC8505].

9.2.4. Perspective of the 6LBR

The 6LBR is unaware that the RPL Root is not the new attachment 6LR of the RUL, so it is not impacted by this specification.

Upon reception of an EDAR message, the 6LBR acts as prescribed by [RFC8505] and returns an EDAC message to the sender.

10. Protocol Operations for Multicast Addresses

Section 12 of [RFC6550] details the RPL support for multicast flows. This support is activated by the MOP of 3 ("Storing Mode of Operation with multicast support") in the DIO messages that form the DODAG. This section also applies if and only if the MOP of the RPLInstance is 3.

The RPL support of multicast is not source-specific and only operates as an extension to the Storing Mode of Operation for unicast packets. Note that it is the RPL model that the multicast packet is passed as a Layer-2 unicast to each of the interested children. This remains true when forwarding between the 6LR and the listener 6LN.

"Multicast Listener Discovery Version 2 (MLDv2) for IPv6" [RFC3810] provides an interface for a listener to register to multicast flows. In the MLD model, the Router is a "querier", and the Host is a multicast listener that registers to the querier to obtain copies of the particular flows it is interested in.

The equivalent of the first Address Registration happens as illustrated in Figure 10. The 6LN, as an MLD listener, sends an unsolicited Report to the 6LR. This enables it to start receiving the flow immediately, and causes the 6LR to inject the multicast route in RPL.

This specification does not change MLD but will operate more efficiently if the asynchronous messages for unsolicited Report and Done are sent by the 6LN as Layer-2 unicast to the 6LR, in particular on wireless.

The 6LR acts as a generic MLD querier and generates a DAO with the Multicast Address as the Target Prefix as described in section 12 of [RFC6550]. As for the Unicast Host routes, the Path Lifetime associated to the Target is mapped from the Query Interval, and set to be larger to account for variable propagation delays to the Root. The Root proxies the MLD exchange as a listener with the 6LBR acting as the querier, so as to get packets from a source external to the RPL domain.

Upon a DAO with a Target option for a multicast address, the RPL Root checks if it is already registered as a listener for that address, and if not, it performs its own unsolicited Report for the multicast address as described in section 5.1 of [RFC3810]. The report is source independent, so there is no Source Address listed.

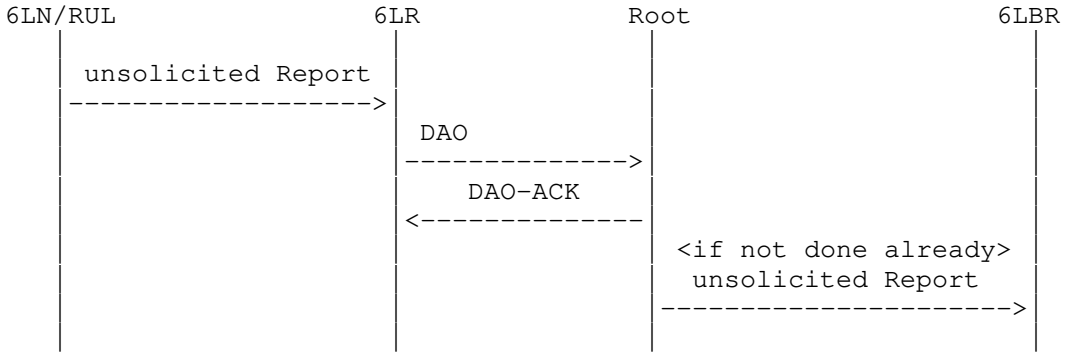

Figure 10: First Multicast Registration Flow

The equivalent of the registration refresh is pulled periodically by the 6LR acting as querier. Upon the timing out of the Query Interval, the 6LR sends a Multicast Address Specific Query to each of its listeners, for each Multicast Address, and gets a Report back that is mapped into a DAO one by one. Optionally, the 6LR MAY send a General Query, where the Multicast Address field is set to zero. In that case, the multicast packet is passed as a Layer-2 unicast to each of the interested children. .

Upon a Report, the 6LR generates a DAO with as many Target Options as there are Multicast Address Records in the Report message, copying the Multicast Address field in the Target Prefix of the RPL Target Option. The DAO message is a Storing Mode DAO, passed to a selection of the 6LR's parents.

Asynchronously to this, a similar procedure happens between the Root and a router such as the 6LBR that serves multicast flows on the Link where the Root is located. Again the Query and Report messages are

source independent. The Root lists exactly once each Multicast Address for which it has at least one active multicast DAO state, copying the multicast address in the DAO state in the Multicast Address field of the Multicast Address Records in the Report message.

This is illustrated in Figure 11:

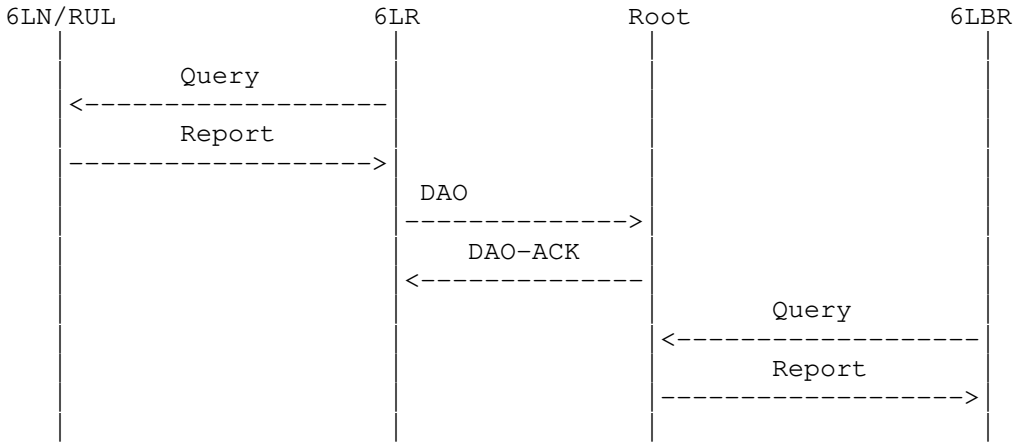

Figure 11: Next Registration Flow

Note that any of the functions 6LR, Root and 6LBR might be collapsed in a single node, in which case the flow above happens internally, and possibly through internal API calls as opposed to messaging.

11. Security Considerations

It is worth noting that with [RFC6550], every node in the LLN is RPL-aware and can inject any RPL-based attack in the network. This specification isolates edge nodes that can only interact with the RPL Routers using 6LoWPAN ND, meaning that they cannot perform RPL insider attacks.

The LLN nodes depend on the 6LBR and the RPL participants for their operation. A trust model must be put in place to ensure that the right devices are acting in these roles, so as to avoid threats such as black-holing, (see [RFC7416] section 7), Denial-Of-Service attacks whereby a rogue 6LR creates a high churn in the RPL network by advertising and removing many forged addresses, or bombing attack whereby an impersonated 6LBR would destroy state in the network by using the status code of 4 ("Removed").

This trust model could be at a minimum based on a Layer-2 Secure joining and the Link-Layer security. This is a generic 6LoWPAN requirement, see Req5.1 in Appendix of [RFC8505].

In a general manner, the Security Considerations in [RFC7416] [RFC6775], and [RFC8505] apply to this specification as well.

The Link-Layer security is needed in particular to prevent Denial-Of-Service attacks whereby a rogue 6LN creates a high churn in the RPL network by constantly registering and deregistering addresses with the R Flag set in the EARO.

[AP-ND] updated 6LoWPAN ND with the called Address-Protected Neighbor Discovery (AP-ND). AP-ND protects the owner of an address against address theft and impersonation attacks in a Low-Power and Lossy Network (LLN). Nodes supporting th extension compute a cryptographic identifier (Crypto-ID), and use it with one or more of their Registered Addresses. The Crypto-ID identifies the owner of the Registered Address and can be used to provide proof of ownership of the Registered Addresses. Once an address is registered with the Crypto-ID and a proof of ownership is provided, only the owner of that address can modify the registration information, thereby enforcing Source Address Validation. [AP-ND] reduces even more the attack perimeter that is available to the edge nodes and its use is suggested in this specification.

Additionally, the trust model could include a role validation to ensure that the node that claims to be a 6LBR or a RPL Root is entitled to do so.

The Opaque field in the EARO enables the RUL to suggest a RPLInstanceID where its traffic is placed. It is also possible for an attacker RUL to include an RPI in the packet. This opens to attacks where a RPL instance would be reserved for critical traffic, e.g., with a specific bandwidth reservation, that the additional traffic generated by a rogue may disrupt. The attack may be alleviated by traditional access control and traffic shaping mechanisms where the 6LR controls the incoming traffic from the 6LN. More importantly, the 6LR is the node that injects the traffic in the RPL domain, so it has the final word on which RPLInstance is to be used for the traffic coming from the RUL, per its own policy.

At the time of this writing, RPL does not have a Route Ownership Validation model whereby it is possible to validate the origin of an address that is injected in a DAO. This specification makes a first step in that direction by allowing the Root to challenge the RUL via the 6LR that serves it.

[EFFICIENT-NPDAO] introduces the ability for a rogue common ancestor node to invalidate a route on behalf of the target node. In this case, the RPL Status in the DCO has the 'A' flag not set, and a NA(EARO) is returned to the 6LN with the R flag not set. This encourages the 6LN to try another 6LR. If a 6LR exists that does not use the rogue common ancestor, then the 6LN will eventually succeed gaining reachability over the RPL network in spite of the rogue node.

12. IANA Considerations

12.1. Fixing the Address Registration Option Flags

Section 9.1 of [RFC8505] creates a Registry for the 8-bit Address Registration Option Flags field. IANA is requested to rename the first column of the table from "ARO Status" to "Bit number".

12.2. Resizing the ARO Status values

Section 12 of [RFC6775] creates the Address Registration Option Status values Registry with a range 0-255.

This specification reduces that range to 0-63, see Section 6.3.

IANA is requested to modify the Address Registration Option Status values Registry so that the upper bound of the unassigned values is 63. This document should be added as a reference. The registration procedure does not change.

12.3. New RPL DODAG Configuration Option Flag

IANA is requested to assign a flag from the "DODAG Configuration Option Flags for MOP 0..6" [USEofRPLinfo] registry as follows:

Bit Number	Capability Description	Reference
1 (suggested)	Root Proxies EDAR/EDAC (P)	THIS RFC

Table 2: New DODAG Configuration Option Flag

12.4. RPL Target Option Registry

This document modifies the "RPL Target Option Flags" registry initially created in Section 20.15 of [RFC6550] . The registry now includes only 4 bits (Section 6.1) and should point to this document as an additional reference. The registration procedure doesn't change.

Section 6.1 also defines a new entry in the Registry as follows:

Bit Number	Capability Description	Reference
1 (suggested)	Advertiser address in Full (F)	THIS RFC

Table 3: RPL Target Option Registry

12.5. New Subregistry for RPL Non-Rejection Status values

This specification creates a new Subregistry for the RPL Non-Rejection Status values for use in the RPL DAO-ACK, DCO, and DCO-ACK messages with the 'A' flag reset, under the RPL registry.

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

* Registration procedure is "IETF Review" [RFC8126].

* Initial allocation is as indicated in Table 4:

Value	Meaning	Reference
0	Unqualified acceptance	THIS RFC / RFC 6550
2..63	Unassigned	

Table 4: Acceptance values of the RPL Status

12.6. New Subregistry for RPL Rejection Status values

This specification creates a new Subregistry for the RPL Rejection Status values for use in the RPL DAO-ACK and DCO messages with the 'A' flag reset, under the RPL registry.

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

* Registration procedure is "IETF Review" [RFC8126].

* Initial allocation is as indicated in Table 5:

Value	Meaning	Reference
0	Unqualified rejection	THIS RFC
1..63	Unassigned	

Table 5: Rejection values of the RPL Status

13. Acknowledgments

The authors wish to thank Ines Robles, Georgios Papadopoulos and especially Rahul Jadhav and Alvaro Retana for their reviews and contributions to this document.

14. 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>.

[RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <https://www.rfc-editor.org/info/rfc3810>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[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>.

[RFC6775] Shelby, Z., Ed., Chakrabarti, S., Nordmark, E., and C. Bormann, "Neighbor Discovery Optimization for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 6775, DOI 10.17487/RFC6775, November 2012, <https://www.rfc-editor.org/info/rfc6775>.

[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>.

[RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <https://www.rfc-editor.org/info/rfc7400>.

[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>.

[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>.

[RFC8504] Chown, T., Loughney, J., and T. Winters, "IPv6 Node Requirements", BCP 220, RFC 8504, DOI 10.17487/RFC8504, January 2019, <https://www.rfc-editor.org/info/rfc8504>.

[RFC8505] Thubert, P., Ed., Nordmark, E., Chakrabarti, S., and C. Perkins, "Registration Extensions for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Neighbor Discovery", RFC 8505, DOI 10.17487/RFC8505, November 2018, <https://www.rfc-editor.org/info/rfc8505>.

[AP-ND] Thubert, P., Sarikaya, B., Sethi, M., and R. Struik, "Address Protected Neighbor Discovery for Low-power and Lossy Networks", Work in Progress, Internet-Draft, draft-ietf-6lo-ap-nd-23, 30 April 2020, <https://tools.ietf.org/html/draft-ietf-6lo-ap-nd-23>.

[USEofRPLinfo]
Robles, I., Richardson, M., and P. Thubert, "Using RPI Option Type, Routing Header for Source Routes and IPv6-in-IPv6 encapsulation in the RPL Data Plane", Work in Progress, Internet-Draft, draft-ietf-roll-useofrplinfo-41, 21 September 2020, <https://tools.ietf.org/html/draft-ietf-roll-useofrplinfo-41>.

[EFFICIENT-NPDAO]
Jadhav, R., Thubert, P., Sahoo, R., and Z. Cao, "Efficient Route Invalidation", Work in Progress, Internet-Draft, draft-ietf-roll-efficient-npdao-18, 15 April 2020, <https://tools.ietf.org/html/draft-ietf-roll-efficient-npdao-18>.

15. Informative References

[RFC4919] Kushalnagar, N., Montenegro, G., and C. Schumacher, "IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals", RFC 4919, DOI 10.17487/RFC4919, August 2007, <https://www.rfc-editor.org/info/rfc4919>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <https://www.rfc-editor.org/info/rfc4862>.

[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>.

[RFC6606] Kim, E., Kaspar, D., Gomez, C., and C. Bormann, "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing", RFC 6606, DOI 10.17487/RFC6606, May 2012, <https://www.rfc-editor.org/info/rfc6606>.

[RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <https://www.rfc-editor.org/info/rfc7039>.

[RFC7228] Bormann, C., Ersue, M., and A. Keranen, "Terminology for Constrained-Node Networks", RFC 7228, DOI 10.17487/RFC7228, May 2014, <https://www.rfc-editor.org/info/rfc7228>.

[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>.

[RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <https://www.rfc-editor.org/info/rfc8415>.

[RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <https://www.rfc-editor.org/info/rfc6282>.

[RFC6687] Tripathi, J., Ed., de Oliveira, J., Ed., and JP. Vasseur, Ed., "Performance Evaluation of the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6687, DOI 10.17487/RFC6687, October 2012, <https://www.rfc-editor.org/info/rfc6687>.

[RFC7416] Tsao, T., Alexander, R., Dohler, M., Daza, V., Lozano, A., and M. Richardson, Ed., "A Security Threat Analysis for the Routing Protocol for Low-Power and Lossy Networks (RPLs)", RFC 7416, DOI 10.17487/RFC7416, January 2015, <https://www.rfc-editor.org/info/rfc7416>.

[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>.

[6BBR] Thubert, P., Perkins, C., and E. Levy-Abegnoli, "IPv6 Backbone Router", Work in Progress, Internet-Draft, draft-ietf-6lo-backbone-router-20, 23 March 2020, <https://tools.ietf.org/html/draft-ietf-6lo-backbone-router-20>.

Appendix A. Example Compression

Figure 12 illustrates the case in Storing Mode where the packet is received from the Internet, then the Root encapsulates the packet to insert the RPI and deliver to the 6LR that is the parent and last hop to the final destination, which is not known to support [RFC8138].

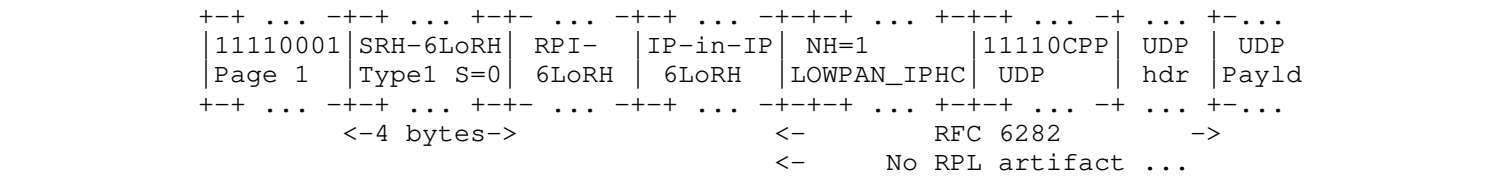

Figure 12: Encapsulation to Parent 6LR in Storing Mode

The difference with the example presented in Figure 19 of [RFC8138] is the addition of a SRH-6LoRH before the RPI-6LoRH to transport the compressed address of the 6LR as the destination address of the outer IPv6 header. In the [RFC8138] example the destination IP of the outer header was elided and was implicitly the same address as the destination of the inner header. Type 1 was arbitrarily chosen, and the size of 0 denotes a single address in the SRH.

In Figure 12, the source of the IP-in-IP encapsulation is the Root, so it is elided in the IP-in-IP 6LoRH. The destination is the parent 6LR of the destination of the inner packet so it cannot be elided. If the DODAG is operated in Storing Mode, it is the single entry in the SRH-6LoRH and the SRH-6LoRH Size is encoded as 0. The SRH-6LoRH is the first 6LoRH in the chain. In this particular example, the 6LR address can be compressed to 2 bytes so a Type of 1 is used. It results that the total length of the SRH-6LoRH is 4 bytes.

In Non-Storing Mode, the encapsulation from the Root would be similar to that represented in Figure 12 with possibly more hops in the SRH-6LoRH and possibly multiple SRH-6LoRHs if the various addresses in the routing header are not compressed to the same format. Note that on the last hop to the parent 6LR, the RH3 is consumed and removed from the compressed form, so the use of Non-Storing Mode vs. Storing Mode is indistinguishable from the packet format.

The SRH-6LoRHs are followed by RPI-6LoRH and then the IP-in-IP 6LoRH. When the IP-in-IP 6LoRH is removed, all the 6LoRH Headers that precede it are also removed. The Paging Dispatch [RFC8025] may also be removed if there was no previous Page change to a Page other than 0 or 1, since the LOWPAN_IPHC is encoded in the same fashion in the default Page 0 and in Page 1. The resulting packet to the destination is the inner packet compressed with [RFC6282].

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
Email: pthubert@cisco.com

Michael C. Richardson
Sandelman Software Works

Email: mcr+ietf@sandelman.ca
URI: http://www.sandelman.ca/

ROLL
Internet-Draft
Intended status: Standards Track
Expires: June 3, 2021

R. Jadhav, Ed.
November 30, 2020

RPL Storing Root-ACK
draft-jadhav-roll-storing-rootack-02

Abstract

This document explains problems with DAO-ACK handling in RPL Storing MOP and provides updates to RFC6550 to solve those problems.

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 June 3, 2021.

Table of Contents

1. Introduction

RPL [RFC6550] specifies a proactive distance-vector routing scheme designed for LLNs (Low Power and Lossy Networks). RPL enables the network to be formed as a DODAG and supports storing mode and non-storing mode of operations. Non-storing mode allows reduced memory resource usage on the nodes by allowing non-BR nodes to operate without managing a routing table and involves use of source routing by the Root to direct the traffic along a specific path. In storing mode of operation the routing happens on hop-by-hop basis and intermediate routers need to maintain routing tables.

DAO messaging helps to install downstream routing paths in the DODAG. DAOs are generated on hop-by-hop basis. DAO may contain multiple RPL Control Options. The Target Option identifies the address prefix for which the route has to be installed and the corresponding Transit Information Option identifies the parameters (such as lifetime, freshness-counter, etc) for the target. The DAO base object contains the 'K' flag indicating that a DAO-ACK is sought by the sender. The DAO, DAO-ACK progresses on hop-by-hop basis all the way till Root. In non-storing MOP, the DAO from the target node is directly addressed to the Root and the Root responds with a DAO-ACK indicating path establishment status. However, in storing MOP, the DAO-ACK is immediately sent by the upstream parent. Thus in case of storing MOP, the target node cannot rely on DAO-ACK as an indication that the end to end (from the target node to Root) path has been established.

This draft highlights various issues with RPL DAO-ACK handling in Storing MOP. Section 4 of [I-D.ietf-roll-rpl-observations] provides

more context to the problem statement. The draft provides requirements to solve the issues and provides an updates to RFC6550 based on these requirements.

1.1. Requirements Language and 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 RFC 2119 [RFC2119].

MOP: Mode of Operation

NS-MOP: RPL Non-Storing Mode of Operation

S-MOP: RPL Storing Mode of Operation

Root-ACK: The Root-ACK syntax is same as DAO-ACK except that the Root-ACK is addressed directly to the peer who owns the target prefix. DAO-ACK in contrast is always sent using link-local IPv6 address in storing MOP.

DelayDAO: Section 9.5 of RFC6550 introduces a delay before the DAO transmission is initiated.

TIO: (Transit Information Option) Section 6.7.8 of RFC6550. TIO is an option usually carried in DAO message and augments control information for the advertised Target.

RUL: (RPL Unaware Leaf) [I-D.ietf-roll-unaware-leaves]

This document uses terminology described in [RFC6550].

2. Problems with DAO-ACK in Storing MOP

Consider the following topology for the subsequent description:

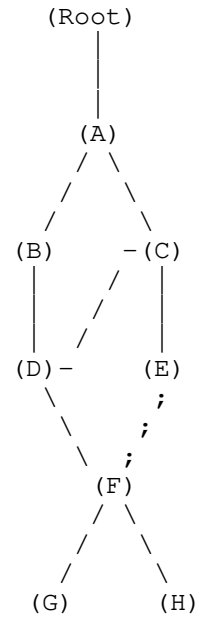

Figure 1: Sample topology

2.1. End to End Path Establishment Indication

Nodes need to know whether the end to end path till the Root has been established before they can initiate application traffic. In case of NS-MOP, the DAO is addressed to the Root from the Target node and the Root sends DAO-ACK directly addressed back to the target node. Thus in case of NS-MOP, the node can make use of this DAO-ACK as an indication whether the necessary routes have been installed. However, in case of Storing MOP, the DAO/DAO-ACK signaling happens at every hop.

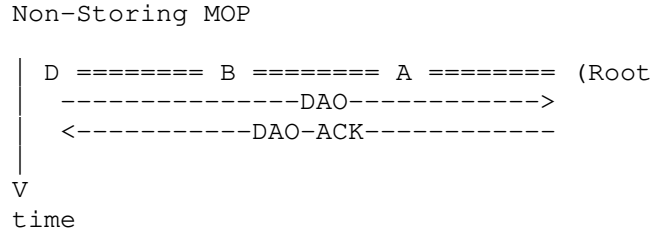

Figure 2: NS-MOP DAO/DAO-ACK handling

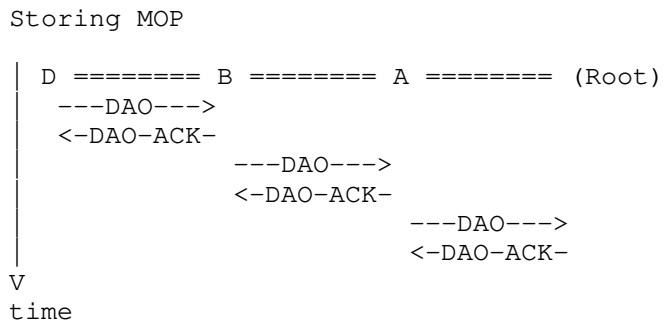

Figure 3: Storing MOP DAO/DAO-ACK handling

Note that in Storing-MOP, the DAO/DAO-ACK signaling happens on hop-by-hop basis and a DelayDAO timer is used before intermediate 6LRs generate the DAO. This would mean that the DAO reaching the Root may take several seconds. The target node should not generate the application traffic unless the end to end path is established.

Consider Figure 1, when node D sends a DAO, the node B receives the DAO and instantly sends back DAO-ACK. Node B then subsequently generates the DAO with Target as Node D and sends it to node A. The DAO with Target as Node D may take time (since the DAO is scheduled with DelayDAO timer by every node) to finally reach the Root at which point the end to end path is established. There is no way for node D to know when the end to end path is established. This information is needed for node D to initiate its application traffic. Initiating application traffic prior to this might almost certainly lead to application packet retries causing congestion in the network.

2.2. Target node is unaware if it needs to retry the DAO

It is possible that the intermediate 6LR goes down while attempting to generate DAO on behalf of the target node. In this case, the target node has no way of knowing to retry the DAO, in which case the route installation may not happen until the target node's DAO lifetime expires.

Consider Figure 1, assume that node A was generating DAO with Target node D and sending it to Root. Node A reboots before attempting to send DAO to Root. Node A has already sent DAO-ACK downstream to node B. In this case, the target node D is not aware that sending DAO has failed somewhere upstream. Note that as per RFC6550 upstream DAO is scheduled based on DelayDAO but DAO-ACK is sent instantaneously on DAO reception from downstream node.

2.3. RPL node acting as router for RULs

An RPL node may act as a router for RPL unware leaves as described in [I-D.ietf-roll-unaware-leaves]. Ideally an RPL node should start accepting RULs solicitation only after making sure that it has established itself in the network first. In Storing-MOP, there is no way to ascertain this.

3. Requirements for Root-ACK handling in Storing MOP

Following are the requirements:

Indicate end to end path establishment The Target node must know when to initiate the application traffic based on end to end path establishment.

Handle multiple targets in DAOs A DAO message may contain multiple Target Options. The Root-ACK mechanism must handle multiple targets in DAO.

Handle DAOs with address prefix RPL DAO Target Option may contain an address prefix i.e., not the full address.

Provide suitable way for target node to retry The Target node must have a way to know and retry the DAO in case the DAO transmission fails enroute.

Backward compatible with current DAO-ACK The current per hop DAO-ACK must function as it is. Legacy nodes should be able to operate without any changes.

4. Root-ACK from Root

The draft defines a way for the RPL Root to send the Root-ACK back directly addressed to the Target node. The Target node can receive the Root-ACK directly thus getting an indication that the end to end path till the Root has been successfully established. The Root-ACK uses the same syntax and message code as DAO-ACK. The only difference is that the Root-ACK is directly addressed to the Target node who owns the advertised prefix in the Target Option.

4.1. Transit Information Option update in DAO message

The Target node indicates that it wishes to receive Root-ACK directly from Root by setting the newly defined 'K' flag in Transit Information Option.

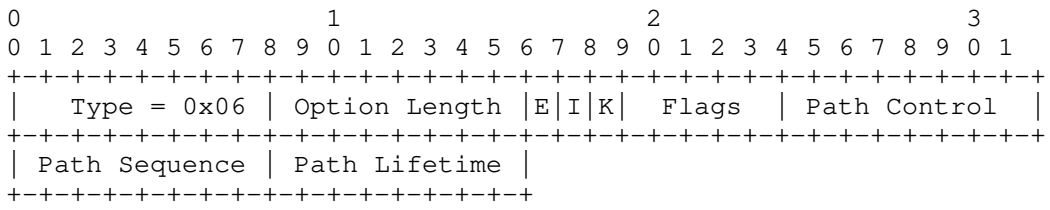

Figure 4: Updated Transit Information Option (New K flag added)

The K flag indicates that the Root of the RPLInstance MUST send a Root-ACK directly to the target node.

4.2. Root sends Root-ACK addressed to Target

On receiving a DAO with Transit Information Option with 'K' flag set, the Root MUST respond with a Root-ACK immediately to the address extracted from the corresponding Target Option.

The Root-ACK MUST contain the Transit Information Option with parameters copied from the DAO's Transit Information Option based on which this Root-ACK was generated. The PathSequence in the Transit Information Option helps the Target node to identify for which DAO it generated it has received the Root-ACK. The DAOSequence in the base Root-ACK(DAO-ACK) base object is ignored by the Target node.

5. IANA Considerations

IANA is requested to allocate bit 2 from the Transit Information Option Flags registry for the 'K' flag (Section 4.1).

6. Security Considerations

This node introduces a new flag in response to which the Root of the DODAG would send a Root-ACK which serves as an indication for the target node that the end to end route/path is established. The Root-ACK indication eventually would be used by the end node for application layer processing such as initiating the application traffic. A malicious node could generate the Root-ACK pre-maturely i.e, before the end-to-end path is established and cause the application to do some processing pre-maturely. However, the application layer would always account for application layer failures and thus shouldn't result in any security issues. This could result in more control overhead which is currently the case where nodes do not support this specification.

A malicious 6LR or 6LN could set the 'K' flag indicating the Root to send a Root-ACK. The Root would generate a Root-ACK for the

indicated target. The Root need not keep any additional state for handling the 'K' flag.

This document assumes that the security mechanisms as defined in [RFC6550] are followed, which means that all the nodes are part of the RPL network because they have the required credentials. A non-secure RPL network needs to take into consideration the risks highlighted in this section as well as those highlighted in [RFC6550].

7. References

7.1. Normative References

[I-D.ietf-roll-unaware-leaves]
Thubert, P. and M. Richardson, "Routing for RPL Leaves", draft-ietf-roll-unaware-leaves-23 (work in progress), November 2020.

[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>.

[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>.

7.2. Informative References

[I-D.ietf-roll-rpl-observations]
Jadhav, R., Sahoo, R., and Y. Wu, "RPL Observations", draft-ietf-roll-rpl-observations-04 (work in progress), May 2020.

Author's Address

Rahul Arvind Jadhav (editor)
Marathahalli
Bangalore, Karnataka 560037
India

Email: rahul.ietf@gmail.com