Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: January 12, 2012

M. Goyal, Ed.
University of Wisconsin
Milwaukee
E. Baccelli
M. Philipp
INRIA
A. Brandt
Sigma Designs
R. Cragie
Gridmerge Ltd
J. Martocci
Johnson Controls

July 11, 2011

Reactive Discovery of Point-to-Point Routes in Low Power and Lossy Networks

draft-ietf-roll-p2p-rpl-04

Abstract

Point to point (P2P) communication between arbitrary IPv6 routers in a Low power and Lossy Network (LLN) is a key requirement for many applications. RPL, the IPv6 Routing Protocol for LLNs, constrains the LLN topology to a Directed Acyclic Graph (DAG) and requires the P2P routing to take place along the DAG links. Such P2P routes may be suboptimal and may lead to traffic congestion near the DAG root. This document specifies a P2P route discovery mechanism, complementary to the RPL base functionality. This mechanism allows an IPv6 router to discover and establish, on demand, a route to another IPv6 router in the LLN such that the discovered route meets specified constraints.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

RPL [I-D.ietf-roll-rpl] provides multipoint-to-point (MP2P) routes from routers in a Low power and Lossy Network (LLN) to a sink by organizing the routers along a Directed Acyclic Graph (DAG) rooted at the sink. The routers determine their position in the DAG so as to optimize their routing cost on the path towards the DAG root. A router advertises its position (the "rank") in the DAG by originating a DODAG Information Object (DIO) message. The DIO message is sent via link-local multicast and also includes information such as the DAG root's identity, routing metrics/constraints [I-D.ietf-roll-routing-metrics] and the objective function (OF) in use. When a router joins the DAG, it determines its own rank in the DAG based on that advertised by its neighbors and originates its own DIO message.

RPL enables point-to-multipoint (P2MP) routing from a router to its descendants in the DAG by allowing a router to send a Destination Advertisement Object (DAO) upwards along the DAG. In non-storing mode operation, a router's DAO contains a list of its preferred DAG parents. The routers unicast their DAOs to the DAG root, which then uses this information to arrive at source-routes from itself to the individual routers. In storing mode operation, a router's DAO carries potentially aggregated information regarding its descendants and other local prefixes reachable through the router. The router sends its DAO to a selected set of DAG parents, which then use this information in their routing tables and in their own DAOs.

RPL also provides mechanisms for point-to-point (P2P) routing between any two routers in the DAG. If the destination is within the source's radio range, the source may directly send packets to the destination. Otherwise, a packet's path from the source to the destination depends on the storing/non-storing operation mode of the DAG. In non-storing mode operation, only the DAG root maintains the "downwards" routing information and hence a packet travels all the way to the DAG root, which then sends it towards its destination using a source route. In storing mode operation, if the destination is a DAG descendant and the source maintains "downwards" hop-by-hop routing state about this descendant, it can forward the packet to a descendant router closer to the destination. Otherwise, the source sends the packet to a DAG parent, which then applies the same set of rules to forward the packet further. Thus, a packet travels up the DAG until it reaches a router that knows of the downwards route to the destination and then it travels down the DAG towards its destination. A router may or may not maintain routing state about a descendant depending on whether its immediate children send it such information in their DAOs. Thus, in the best case with storing mode operation, the "upwards" segment of the P2P route between a source

and a destination ends at the first common ancestor of the source and the destination. In the worst case, the "upwards" segment would extend all the way to the DAG root. In both storing and non-storing mode operations, if the destination did not originate a DAO, the packet will travel all the way to the DAG's root, where it will be dropped.

The P2P routing functionality available in RPL may be inadequate for applications in the home and commercial building domains for the following reasons [I-D.brandt-roll-rpl-applicability-home-building] [RFC5826][RFC5867]:

- o The need to maintain routes "proactively", i.e., every possible destination in the DAG must originate a DAO.

- o Depending on the network topology and OF/metrics in use, the constraint to route only along a DAG may cause significantly suboptimal P2P routes and severe traffic congestion near the DAG root.

Thus, there is a need for a mechanism that provides source-initiated discovery of a P2P route that need not be along an existing DAG. This document describes such a mechanism, complementary to the basic RPL functionality.

The specified mechanism is based on a reactive on-demand approach, which enables a router to discover a route to another router in the LLN without any restrictions regarding the existing DAG-membership of the links that the route may use. The specified mechanism allows for the discovery of sources routes as well as hop-by-hop ones. The discovered route may not be the best available but is guaranteed to satisfy the desired constraints in terms of the routing metrics and is thus considered "good enough" from the application's perspective.

A complementary functionality, necessary to help decide whether to initiate a route discovery, is a mechanism to measure the end-to-end cost of an existing route. Section 4 provides further details on how such functionality, described in [I-D.ietf-roll-p2p-measurement], can be used to determine the metric constraints for use in the route discovery mechanism described in this document.

2. The Use Cases

The mechanisms described in this document are intended to be employed as complementary to RPL in specific scenarios that need point-to-point (P2P) routes between arbitrary routers.

One use case, common in a home environment, involves a remote control (or a motion sensor) that suddenly needs to communicate with a lamp module, whose network address is a-priori known. In this case, the source of data (the remote control or the motion sensor) must be able to discover a route to the destination (the lamp module) "on demand".

Another use case, common in a large commercial building environment, involves a large LLN deployment where P2P communication along a particular DAG among hundreds (or thousands) of routers creates severe traffic congestion near that DAG's root, and thus routes across this DAG are desirable.

The use cases also include scenarios where energy or latency constraints are not satisfied by the P2P routes along a DAG because they involve traversing many more intermediate routers than necessary to reach the destination.

3. Terminology

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 [RFC2119].

Additionally, this document uses terminology from [I-D.ietf-roll-terminology] and [I-D.ietf-roll-rpl]. This document introduces the following terms:

Origin : The RPL node initiating the route discovery.

Target : The RPL node at the other end point of the route(s) to be discovered.

Intermediate Router: An RPL router that is neither the origin nor the target.

Forward Route: A route in the forward direction, i.e., from the origin to the target.

Backward Route: A route in the backward direction, i.e., from the target to the origin.

Bidirectional Route: A route that can be used in both forward and backward directions.

Source Route: A complete and ordered list of routers that can be used by a packet to travel from a source to a destination node.

Hop-by-hop Route: The route characterized by each router on the route using its routing table to determine the next hop on the route.

4. Applicability

The route discovery mechanism, described in this document, may be invoked by an origin when no route exists between itself and the target or when the existing routes do not satisfy the desired performance requirements. The mechanism is designed to discover and establish one hop-by-hop route or discover one or more source routes such that the discovered route(s) meet the specified constraints. In some application contexts, the constraints that the discovered route(s) must satisfy are intrinsically known or can be specified by the application. For example, an origin that expects a target to be less than 5 hops away may use "hop-count < 5" as the constraint. In other application contexts, the origin may need to measure the cost of an existing route to the target to determine the constraints. For example, an origin that measures the total ETX of its along-DAG route to the target to be 20 may use "ETX < x*20", where x is a fraction that the origin decides, as the constraint. The functionality required to measure the cost of an existing route between the origin and the target is described in [I-D.ietf-roll-p2p-measurement]. In case, there is no existing route between the origin and target or the cost measurement for the existing route fails, the origin will have to guess the constraints used in the initial route discovery. Once, the initial route discovery succeeds or fails, the origin will have a better estimate for the constraints to be used in the subsequent route discovery.

This document describes an on-demand discovery mechanism for P2P routes that is complementary to the proactive routes offered by RPL base functionality. The mechanism described in this document may result in discovery of better P2P routes than the ones available along a DAG designed to optimize routing cost to the DAG's root. The improvement in route quality depends on a number of factors including the network topology, the routing metrics in use and the prevalent conditions in the network. A network designer may take in consideration both the benefits (potentially better routes; no need to maintain routes proactively) and costs (control messages generated during the route discovery process) when using this mechanism.

5. Functional Overview

This section contains a high level description of the route discovery mechanism proposed in this document.

The P2P route discovery takes place by forming a temporary DAG rooted at the origin. The DIOs used to create the temporary DAG also carry the following information:

o The target

o The relevant routing metrics

o The constraints that the discovered route must satisfy. These constraints also limit how far the Discovery message may travel.

o The nature of the route(s) to be discovered: hop-by-hop or source routes. This specification allows for the discovery of one hop-by-hop route or up to four source routes in the forward direction.

o The desired number of routes (if source routes are being discovered)

o Whether the route(s) need to be bidirectional. If bidirectional route(s) are being discovered, the target may store the route in backward direction for use as a source route. This specification does not provide for the establishment of backward hop-by-hop routes.

As the routers join the temporary DAG, they keep track of the best (partial) route(s) they have seen and advertise these routes, along with the corresponding routing metrics, in their DIOs. The routing metrics are measured in forward direction unless bidirectional routes are being discovered, in which case the measurement of routing metrics need to take in account both forward and backward directions. A router, including the target, discards a received DIO if the aggregated routing metrics on the route advertised by the DIO do not satisfy the listed constraints. These constraints can be used to limit the propagation of DIO messages used for P2P route discovery. A router may also discard a received DIO if it does not wish to be a part of the discovered route due to limited resources or due to policy reasons.

When the target receives a DIO, it checks whether the route advertised therein satisfies the routing constraints. If yes, the target may select the route for further processing as described next. This document does not specify a particular method for the target to select a route among the ones that satisfy the route constraints. Example selection methods include selecting any route that meets the constraints or selecting the best route(s) discovered over a certain time period.

If one or more source routes are being discovered, the target sends the discovered source routes to the origin via Discovery Reply Object (DRO) messages, defined in Section 7, with one DRO message carrying one discovered route. On receiving a DRO message, the origin stores the route contained therein in its memory.

If a hop-by-hop route is being discovered, the target sends a DRO message to the origin after selecting a suitable route among the ones that satisfy the route constraints. The DRO message travels towards the origin along the discovered route, establishing state for this route in the routers on the path.

The target may store a discovered route in its memory if it is bidirectional and use it as a backward source-route to send packets to the origin.

The target may request the origin to acknowledge the receipt of a DRO message by sending back a DRO Acknowledgement (DRO-ACK) message (Section 8). The origin unicasts a DRO-ACK message to the target. When the target does not receive the requested DRO-ACK within a certain time interval of sending a DRO, it resends the DRO message carrying the same route as before.

6. P2P Route Discovery By Creating a Temporary DAG

RPL uses DIO message propagation to build a DAG. The DIO message travels via IPv6 link-local multicast. Each router joining the DAG determines a rank for itself and ignores the subsequent DIO messages received from lower (higher in numerical value) ranked neighbors. Thus, the DIO messages propagate outward from the DAG root rather than return inward towards the DAG root. The DIO message generation at a router is further controlled by a Trickle timer that allows a router to avoid generating unnecessary messages [RFC6206]. The link-local multicast based propagation, Trickle-controlled generation and the rank-based poisoning of messages traveling in the wrong direction (towards the DAG root) provide powerful incentives to use the DIO message for P2P route discovery by creating a "temporary" DAG. Such an approach also allows the reuse of the routing metrics, objective function and packet forwarding framework developed for RPL. This document defines a new RPL option, Route Discovery Option (RDO), which when carried inside a DIO message identifies that message as doing P2P route discovery by creating a temporary DAG as specified in this document.

The use of trickle timers to delay the propagation of DIO messages may cause some nodes to generate these messages even when the desired routes have already been discovered. In order to preempt the generation of such unnecessary messages, the target may set a "stop" bit in the DRO message to let the nodes in the LLN know about the completion of the route discovery process.

6.1. The Route Discovery Option

```
 0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |   Type = 9    | Option Length |D|H| N | Compr | L |     Rem    |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 |                            Target                             |
 |                                                               |
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 |                        Address[1..n]                          |
 |                                                               |
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 1: Format of the Route Discovery Option

In order to perform P2P route discovery as specified in this document, a DIO MUST carry a Route Discovery Option (RDO) illustrated in Figure 1. A Discovery Reply Object (DRO) message, defined in Section 7, MUST also carry a Route Discovery Option. A DIO/DRO message MUST NOT carry more than one Route Discovery Option. A router MUST discard a DIO/DRO if it contains more than one Route Discovery Option. A Route Discovery Option consists of the following fields:

- Option Type: 0x09 (to be confirmed by IANA).

- Option Length: The length of the option in bytes.

- Direction (D): This flag indicates the direction in which the desired routes should be optimized. The flag is set to 1 if the routes are to be optimized for use in both forward and backward directions. If the discovered routes need be optimized in the forward direction only, the flag is reset to 0. Note that the discovered routes must have bidirectional reachability irrespective of the value of D flag. This is because the DRO messages travel from the target to the origin along one of the discovered routes. When the Route Discovery Option is carried inside a DIO, the link-level metric objects contained in the DIO SHOULD be measured in the direction indicated by the D flag. When

the Route Discovery Option is carried inside a DRO message, this flag should be set to zero on transmission and ignored on reception.

o Hop-by-hop (H): This flag is set to 1 if a hop-by-hop route is desired. The flag is reset to zero if source routes are desired. This specification allows for the establishment of one hop-by-hop route and up to four source routes in the forward direction. This specification does not allow for the establishment of hop-by-hop routes in the backward direction. If a bidirectional route is discovered, the target MAY use the route in backward direction as a source route to reach the origin, irrespective of the value of H flag.

o Number of Routes (N): When the Route Discovery Option is being carried inside a DIO and the source routes are being discovered, the value in this field plus one indicates the desired number of routes. When a hop-by-hop route is being discovered or when the Route Discovery Option is being carried inside a DRO message, this field MUST be set to zero on transmission and ignored on reception.

o Compr: 4-bit unsigned integer indicating the number of prefix octets that are elided from the Target field and the Address vector. For example, Compr value will be 0 if full IPv6 addresses are carried in the Target field and the Address vector.

o Life Time (L): A 2-bit field that indicates the suggested life time of the temporary DAG, i.e., the suggested duration a router joining the temporary DAG must maintain its membership in the DAG. The mapping between the values in this field and the minimum life time of the temporary DAG is as follows:

* 0x00: 1 second;

* 0x01: 4 seconds;

* 0x02: 16 seconds;

* 0x03: 64 seconds;

Note that a router MAY detach from the temporary DAG sooner if it receives a DRO message concerning this DAG with "stop" bit set. When a Route Discovery Option is included inside a DRO message, this field MUST be set to zero on transmission and ignored on reception.

o Rem: When the Route Discovery Option is carried inside a DIO, this field indicates the number of empty fields inside the Address vector. When the Route Discovery Option is carried inside a DRO message, this field indicates the number of fields in the Address vector yet to be visited.

o Target: The IPv6 address of the target after eliding Compr number of prefix octets.

o Address[1..n]: A vector of IPv6 addresses representing a (partial) route in the forward direction:

  * Each element in the vector has size (16 - Compr) octets.

  * The total number of elements inside the Address vector is given by n = (Option Length - 4 - (16 - Compr))/(16 - Compr).

  * When the Route Discovery Option is carried inside a DIO, the Address vector is used to accumulate a route optimized in the direction specified by the D field.

  * The IPv6 addresses in the Address vector MUST be accessible in both forward and backward directions. Accessibility in the backward direction is required because the DRO message uses the route accumulated in the Address vector to travel from the target to the origin.

  * The Address vector MUST carry the accumulated route in the forward direction, i.e., the first element in the Address

vector must contain the IPv6 address of the router next to the origin and so on.

* The origin and target addresses MUST NOT be included in the Address vector.

* A router adding its address to the vector MUST ensure that its address does not already exist in the vector. A router specifying a complete route in the Address vector MUST ensure that the vector does not contain any address more than once.

* The Address vector MUST NOT contain any multicast addresses.

* When the Route Discovery Option is carried inside a DRO message, the Address vector MUST contain a complete route between the origin and the target such that the first element in the vector contains the IPv6 address of the router next to the origin and the last element contains the IPv6 address of the router next to the target.

6.2. Setting a DIO Carrying a Route Discovery Option

The Base Object in a DIO message carrying a Route Discovery Option MUST be set in the following manner:

o RPLInstanceID: RPLInstanceID MUST be a local value as described in Section 5.1 of [I-D.ietf-roll-rpl]. The origin MUST NOT use the same RPLInstanceID in two or more concurrent route discoveries. The origin MAY use the same RPLInstanceID value to establish hop-by-hop P2P routes to different target routers.

o Version Number: MUST be set to zero. The temporary DAG used for P2P route discovery does not exist long enough to have new versions.

o Grounded (G) Flag: MUST be cleared since this DAG is temporary in nature and MUST NOT be used for routing purpose.

o Mode of Operation (MOP), DTSN: These fields MUST be set to value 0 since this DAG does not support downward routing.

o DODAGPreference (Prf): This field MUST be set to value 0 (least preferred).

o DODAGID: This field MUST be set to the IPv6 address of the origin.

o The other fields in the Base Object can be set in the desired fashion as per the rules described in [I-D.ietf-roll-rpl].

The DODAG Configuration option, carried in the DIO message, MUST be set in the following manner:

o MaxRankIncrease: This field MUST be set to 0 to disable local repair of the temporary DAG.

o Trickle parameters SHOULD be set as described in Section 6.4.

o The Default Lifetime and Lifetime Unit parameters in DODAG Configuration option indicate the life time of the state the routers maintain for a hop-by-hop route established using the mechanism described in this draft.

o The other fields in the DODAG Configuration option, including the OCP, can be set in the desired fashion as per the rules described in [I-D.ietf-roll-rpl].

A DIO, carrying a Route Discovery Option, MUST NOT carry any Route Information or Prefix Information options described in

[I-D.ietf-roll-rpl].

6.3. Joining a Temporary DAG

When a router joins a temporary DAG advertized by a DIO carrying a Route Discovery Option, it SHOULD maintain its membership in the DAG for the suggested Life Time duration listed in the Route Discovery Option. Maintaining membership in the DAG implies remembering:

o The RPLInstanceID, the DODAGID and the DODAGVersionNumber for the temporary DAG;

o The router's rank in the temporary DAG;

o The best values of the routing metrics, along with the associated route(s) from the origin until this router (carried inside the Route Discovery Option) in the DIOs received so far.

The only purpose of a temporary DAG's existence is to facilitate the propagation of the Discovery messages. The temporary DAG MUST NOT be used to route packets. A router SHOULD detach from the temporary DAG once the duration of its membership in the DAG has exceeded the DAG's suggested life time. A router MAY detach from a temporary DAG sooner when it receives a DRO about the temporary DAG with stop flag set.

6.4. Trickle Operation For DIOs Carrying a Route Discovery Option

An RPL router uses a Trickle timer [RFC6206] to control DIO transmissions. The Trickle control of DIO transmissions provides quick resolution of any "inconsistency" while avoiding redundant DIO transmissions. The Trickle algorithm also imparts protection against loss of DIOs due to inherent lack of reliability in wireless communication. When controlling the transmissions of a DIO carrying a Route Discovery Option, a Trickle timer SHOULD follow the following rules:

- The receipt of a DIO, that allows the router to advertise a better route (in terms of the routing metrics and the OCP in use) than before, is considered "inconsistent" and hence resets the Trickle timer. Note that the first receipt of a DIO advertising a particular temporary DAG is always considered an inconsistent event under this rule.

- The receipt of a DIO, that advertises a better route than the router but does not lead to the router advertising a better route itself, is considered "consistent".

- The receipt of a DIO, that advertises as good a route as the router itself, is considered "consistent".

- The receipt of a DIO, that advertises a worse route than what the router advertises, is considered neither "consistent" nor "inconsistent", i.e., the receipt of such a DIO has no impact on the Trickle operation.

o The recommended values of Imin and Imax are same as in base RPL specification [I-D.ietf-roll-rpl], i.e., 8ms and 2.3 hours respectively.

o The recommended value of redundancy constant "k" is 1. With this value of "k", a DIO transmission will be suppressed if the router receives even a single "consistent" DIO during a timer interval.

6.5. Processing a DIO Carrying a Route Discovery Option

The rules for DIO processing and transmission, described in Section 8 of RPL [I-D.ietf-roll-rpl], apply to DIOs carrying a Route Discovery option as well except as modified in this document.

The following rules for processing a DIO carrying a Route Discovery Option apply to both intermediate routers and the target.

A router SHOULD discard a received DIO with no further processing if it does not have bidirectional reachability with the neighbor that originated the received DIO. This is to ensure that a discovered route can be used to send a DRO message from the target to the origin. Note that bidirectional reachability does not mean that the link must have the same values for a routing metric in both directions. A router SHOULD update the values of the link-level routing metrics included inside the DIO in the direction indicated by the D flag in the Route Discovery Option. If the D flag is 0, i.e., the discovered routes need not be bidirectional, the link-level routing metrics SHOULD be measured in the forward direction, i.e., towards the node receiving the DIO. If the D flag is 1, i.e., bidirectional routes are desired, the link-level routing metrics SHOULD be calculated so as to take in account the metric's value in both forward and backward directions.

A router MUST discard the DIO with no further processing if it can not evaluate the mandatory route constraints listed in the DIO or if the routing metric values do not satisfy one or more of the mandatory constraints.

6.6. Additional Processing of a DIO Carrying a Route Discovery Option

At An Intermediate Router

When an intermediate router receives a DIO containing a Route Discovery Option, it MUST determine whether this DIO advertises a better route than the router itself and whether the receipt of the DIO would allow the router to advertise a better route than before. Accordingly, the router SHOULD consider this DIO as consistent/ inconsistent from Trickle perspective as described in Section 6.4. If the received DIO would allow the router to improve the route it advertises, the router MUST add its IPv6 address to the route inside the received DIO at location Address[n-Rem+1] and store this route in memory for inclusion in its future DIOs. When an intermediate router adds itself to a route, it MUST ensure that the IPv6 address added to the route is accessible in both forward and backward directions. To improve the diversity of the routes being discovered, an intermediate router SHOULD remember multiple partial routes, the best it knows in terms of the routing metrics, that it can advertise in the Route Discovery Option inside its DIO. When the router generates its DIO, it SHOULD randomly select the partial route to be included in the Route Discovery Option from the set of best routes it has seen so far.

6.7. Additional Processing of a DIO Carrying a Route Discovery Option At The Target

The target discards a received DIO with no further processing if the routing metrics inside the DIO do not satisfy the the mandatory constraints. Otherwise, the target MAY select the route contained in the Route Discovery Option for further processing. This document does not prescribe a particular method for the target to select such routes. Example selection methods include selecting the desired number of routes as they are identified or selecting the best routes discovered over a certain time period. If multiple routes are desired, the target SHOULD avoid selecting routes that have large segments in common. If a discovered route is bidirectional (D=1), the target MAY store the route in backward direction, obtained by reversing the discovered forward route, for use as a source route to reach the origin. After selecting a route, the target sends a Discovery Reply Object (DRO) message back to the origin (identified by the DODAGID field in the DIO). In this DRO, the target includes a Route Discovery Option that contains the selected route inside the Address vector. The Route Discovery Option included in the DRO message MUST copy the H flag from the Route Discovery Option inside the received DIO message. The other fields inside the Route Discovery Option MUST be set as specified in Section 6.1. The mechanism for the propagation of DRO messages is described in Section 7.

The target MAY set the A flag inside the DRO message if it desires the origin to send back a DRO-ACK message on receiving the DRO. In this case, the target waits for DRO_ACK_WAIT_TIME duration for the DRO-ACK message to arrive. Failure to receive the DRO-ACK message within this time duration causes the target to retransmit the DRO message. The target MAY retransmit the DRO message in this fashion up to MAX_DRO_RETRANSMISSIONS times.

The target MAY include a Metric Container Option in the DRO message. This Metric Container contains the end-to-end routing metric values for the route specified in the Route Discovery Option. The target MAY set the stop flag inside the DRO message if it has already selected the desired number of routes. A target MUST NOT forward a DIO carrying a Route Discovery option any further.

7. The Discovery Reply Object (DRO)

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| RPLInstanceID |    Version    |S|A|Seq|      Reserved         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                            DODAGID                            |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Option(s)...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+...
```

Figure 2: Format of the Discovery Reply Object (DRO)

This document defines a new RPL Control Message type, the Discovery Reply Object (DRO) with code 0x04 (to be confirmed by IANA), that serves one of the following functions:

- Carry a discovered source route from the target to the origin;

- Establish a hop-by-hop route as it travels from the target to the origin.

A DRO message MAY also serve the function of letting the routers in the LLN know that a P2P route discovery is complete and no more DIO messages need to be generated for the corresponding temporary DAG. A

DRO message MUST carry one Route Discovery Option and travel from the target to the origin via link-local multicast along the route

specified in the Route Discovery Option.

The format for a Discovery Reply Object (DRO) is shown in Figure 2. A DRO consists of the following fields:

- o RPLInstanceID: The RPLInstanceID of the temporary DAG used for route discovery.

- o Version: The Version of the temporary DAG used for route discovery.

- o Stop (S): This flag, when set by the target, indicates that the P2P route discovery is over. The routers, receiving such a DRO, SHOULD cancel any pending DIO transmissions for the temporary DAG created for the route discovery and MAY detach from this DAG immediately. Note that the stop flag serves to stop further DIO transmissions for a P2P route discovery but it does not affect the processing of DRO messages at either the origin or the intermediate routers. In other words, a router (the origin or an intermediate router) MUST continue to process the DRO messages even if an earlier DRO message (with same RPLInstanceID, DODAGID and Version Number fields) had the stop flag set.

- o Ack Required (A): This flag, when set by the target, indicates that the origin SHOULD unicast a DRO-ACK message to the target when it receives the DRO.

- o Sequence Number (Seq): This 2-bit field indicates the sequence number for the DRO. This field is relevant when the A flag is set, i.e., the target requests an acknowledgement from the origin for a received DRO. The origin includes the RPLInstanceID, the DODAGID and the Sequence Number of the received DRO inside the DRO-ACK message it sends back to the target.

- o Reserved: These bits are reserved for future use. These bits MUST be set to zero on transmission and MUST be ignored on reception.

- o DODAGID: The DODAGID of the temporary DAG used for route discovery. The DODAGID also identifies the origin. The

RPLInstanceID, the Version and the DODAGID together uniquely identify the temporary DAG used for route discovery and can be copied from the DIO message advertizing the temporary DAG.

- o Options: The DRO message MUST carry one Route Discovery Option that MUST specify a complete route between the target and the origin. The DRO message MAY carry a Metric Container Option that contains the aggregated routing metrics values for the route specified in Route Discovery Option.

7.1. Processing a DRO At An Intermediate Router

When a router receives a DRO message that does not list its IPv6 address in the DODAGID field, the router MUST process the received message in the following manner:

- o If the stop flag inside the received DRO is set and the router currently belongs to the temporary DAG identified by the (RPLInstanceID, DODAGID and Version fields of the) DRO, the router SHOULD cancel any pending DIO transmissions for this temporary DAG. Additionally, the router MAY detach from the temporary DAG immediately.

- o An intermediate router MUST ignore any Metric Container Option contained in the DRO message.

- o If Address[Rem] element inside the Route Discovery Option lists the router's own IPv6 address, the router is a part of the route carried in the Route Discovery Option. In this case, the router MUST do the following:

  * If the H flag inside the Route Discovery Option inside the DRO message is set, the router SHOULD store the state for the forward hop-by-hop route carried inside the Route Discovery Option. This state consists of:

    + The RPLInstanceID and the DODAGID fields of the DRO.

    + The route's destination, the target (identified by Target field in Route Discovery Option).

    + The IPv6 address of the next hop, Address[Rem+1] (unless Rem value equals the number of elements in the Address vector, in which case the target itself is the next hop).

The router MUST drop the DRO message without further processing if the H flag inside the Route Discovery Option is set but the router chooses not to store the state for the hop-by-hop route.

* If the router already maintains a hop-by-hop state listing the target as the destination and carrying same RPLInstanceID and DODAGID fields as the received DRO and the next hop information in the state does not match the next hop indicated in the received DRO, the router MUST drop the DRO message with no further processing.

* The router MUST decrement the Rem field inside the Route Discovery Option and send the DRO further via link-local

multicast.

7.2. Processing a DRO At The Origin

When a router receives a DRO message that lists its IPv6 address in the DODAGID field, the router recognizes itself as the origin for the corresponding P2P route discovery and processes the Route Discovery Option contained in the DRO in the following manner.

If the stop flag inside the received DRO is set and the origin still belongs to the temporary DAG it initiated, it SHOULD cancel any pending DIO transmissions for this temporary DAG. Additionally, the origin MAY detach from the temporary DAG immediately.

If the Route Discovery Option inside the DRO identifies the discovered route as a source route (H=0), the origin SHOULD store in its memory the discovered route contained in the Address vector.

If the Route Discovery Option inside the DRO identifies the discovered route as a hop-by-hop route (H=1), the origin SHOULD store in its memory the state for the discovered route in the manner described in Section 7.1.

If the received DRO message contains a Metric Container Option as well, the origin MAY store the values of the routing metrics associated with the discovered route in its memory. This information may be useful in formulating the constraints for any future P2P route discovery to the target.

If the A flag is set to one in the received DRO message, the origin SHOULD generate a DRO-ACK message as described in Section 8 and unicast the message to the target. The origin MAY source route the DRO-ACK message to the target using the route contained in the received DRO. If the received DRO established a hop-by-hop route to the target, the origin MAY send the DRO-ACK message along this route. Section 9 describes how a packet may be forwarded along a route discovered using the mechanism described in this document.

8. The Discovery Reply Object Acknowledgement (DRO-ACK)

```
     0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     | RPLInstanceID |     Version     |Seq|         Reserved          |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
     |                                                               |
     |                            DODAGID                            |
     |                                                               |
     |                                                               |
     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: Format of the Discovery Reply Object Acknowledgement (DRO-ACK)

A DRO message may fail to reach the origin due to a number of reasons. Unlike the DIO messages that benefit from Trickle-controlled retransmissions, the DRO messages are prone to loss due to reasons associated with wireless communication. Since a DRO message travels via link-local multicast, it can not use link-level acknowledgements to improve the reliability of its transmission. Also, an intermediate router may drop the DRO message (e.g., because of its inability to store the state for the hop-by-hop route the DRO is establishing). To protect against the potential failure of a DRO message to reach the origin, the target MAY request the origin to send back a DRO Acknowledgement (DRO-ACK) message on receiving a DRO message. Failure to receive such an acknowledgement within the DRO_ACK_WAIT_TIME interval of sending the DRO message forces the target to resend the message.

DRO Acknowledgement is a new RPL Control Message type with code 0x05 (to be confirmed by IANA). A DRO-ACK message MUST travel as a unicast message from the origin to the target. The format for a DRO-ACK message is shown in Figure 3. Various fields in a DRO-ACK message MUST have the same values as the corresponding fields in the DRO message. The field marked as "Reserved" MUST be set to zero on transmission and MUST be ignored on reception.

9. Packet Forwarding Along a P2P Route

This document specifies a mechanism to discover P2P routes, which can be either source routes or hop-by-hop ones. A packet MAY use an RH4 header [I-D.ietf-6man-rpl-routing-header] to travel along a P2P source route. Travel along a P2P hop-by-hop route requires specifying the RPLInstanceID and the DODAGID to identify the route.

This is because P2P route discovery does not use globally unique RPLInstanceID values and hence both the RPLInstanceID, which is a local value assigned by the origin, and the DODAGID, which is an IPv6 address belonging to the origin, are required to uniquely identify a P2P hop-by-hop route to a particular destination. A packet MAY include an RPL option [I-D.ietf-6man-rpl-option] inside the IPv6 hop-by-hop options header to travel along a P2P hop-by-hop route. In this case, the origin MUST set the DODAGID of the P2P route as the source IPv6 address of the packet. Further, the origin MUST specify the RPLInstanceID, associated with the P2P route, inside the RPL option and set the O flag inside the RPL option to 1. A router receiving this packet will check the O flag inside the RPL option and correctly infer the source IPv6 address of the packet as the DODAGID of the hop-by-hop route to be used for forwarding the packet further.

10. Security Considerations

TBA

11. IANA Considerations

TBA

12. Acknowledgements

Authors gratefully acknowledge the contributions of the following individuals (in alphabetical order) in the development of this document: Dominique Barthel, Thomas Clausen, Richard Kelsey, Phil Levis, Zach Shelby, Pascal Thubert and JP Vasseur.

13. References

13.1. Normative References

[I-D.ietf-roll-p2p-measurement]
Goyal, M., Baccelli, E., Brandt, A., Cragie, R., Martocci, J., and C. Perkins, "A Mechanism to Measure the Quality of a Point-to-point Route in a Low Power and Lossy Network", draft-ietf-roll-p2p-measurement-00 (work in progress), April 2011.

[I-D.ietf-roll-routing-metrics]
Vasseur, J., Kim, M., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics used for Path Calculation in Low

Power and Lossy Networks", draft-ietf-roll-routing-metrics-19 (work in progress), March 2011.

[I-D.ietf-roll-rpl]
Winter, T., Thubert, P., Brandt, A., Clausen, T., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., and J. Vasseur, "RPL: IPv6 Routing Protocol for Low power and Lossy Networks", draft-ietf-roll-rpl-19 (work in progress), March 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, March 2011.

13.2. Informative References

[I-D.brandt-roll-rpl-applicability-home-building]
Brandt, A., Baccelli, E., and R. Cragie, "Applicability Statement: The use of RPL in Building and Home Environments", draft-brandt-roll-rpl-applicability-home-building-01 (work in progress), November 2010.

[I-D.ietf-6man-rpl-option]
Hui, J. and J. Vasseur, "RPL Option for Carrying RPL Information in Data-Plane Datagrams", draft-ietf-6man-rpl-option-03 (work in progress), March 2011.

[I-D.ietf-6man-rpl-routing-header]
Hui, J., Vasseur, J., Culler, D., and V. Manral, "An IPv6 Routing Header for Source Routes with RPL", draft-ietf-6man-rpl-routing-header-03 (work in progress), March 2011.

[I-D.ietf-roll-terminology]
Vasseur, J., "Terminology in Low power And Lossy Networks", draft-ietf-roll-terminology-05 (work in progress), March 2011.

[RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5826, April 2010.

[RFC5867] Martocci, J., De Mil, P., Riou, N., and W. Vermeylen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5867, June 2010.

Authors' Addresses

Mukul Goyal (editor)
University of Wisconsin Milwaukee
3200 N Cramer St
Milwaukee, WI 53201
USA

Phone: +1 414 2295001
Email: mukul@uwm.edu

Emmanuel Baccelli
INRIA

Phone: +33-169-335-511
Email: Emmanuel.Baccelli@inria.fr
URI: http://www.emmanuelbaccelli.org/

Matthias Philipp
INRIA

Email: Matthias.Philipp@inria.fr

Anders Brandt
Sigma Designs
Emdrupvej 26A, 1.
Copenhagen, Dk-2100
Denmark

Phone: +45-29609501
Email: abr@sdesigns.dk

Robert Cragie
Gridmerge Ltd
89 Greenfield Crescent
Wakefield WF4 4WA
UK

Phone: +44-1924910888
Email: robert.cragie@gridmerge.com

Jerald Martocci
Johnson Controls
507 E Michigan St
Milwaukee, WI 53202
USA

Phone: +1 414-524-4010
Email: jerald.p.martocci@jci.com