

Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: September 21, 2013

M. Goyal, Ed.
University of Wisconsin
Milwaukee
E. Baccelli
M. Philipp
INRIA
A. Brandt
Sigma Designs
J. Martocci
Johnson Controls
March 20, 2013

Reactive Discovery of Point-to-Point Routes in Low Power and Lossy
Networks
draft-ietf-roll-p2p-rpl-17

Abstract

This document specifies a point-to-point route discovery mechanism, complementary to the RPL core functionality. This mechanism allows an IPv6 router to discover "on demand" routes to one or more IPv6 routers in the LLN such that the discovered routes meet specified metrics constraints.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 21, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal

Internet-Draft

[draft-ietf-roll-p2p-rpl-17](#)

March 2013

Provisions Relating to IETF Documents

(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	4
1.1.	Known Issues/Future Work	4
2.	The Use Cases	5
3.	Terminology	5
4.	Applicability	6
5.	Functional Overview	8
6.	P2P Route Discovery Mode Of Operation	10
6.1.	Setting a P2P Mode DIO	11
7.	P2P Route Discovery Option (P2P-RD0)	14
8.	The P2P Discovery Reply Object (P2P-DR0)	18
8.1.	Secure P2P-DR0	20
8.2.	Setting a P2P-RD0 Carried in a P2P Discovery Reply Object	20
9.	P2P-RPL Route Discovery By Creating a Temporary DAG	21
9.1.	Joining a Temporary DAG	21
9.2.	Trickle Operation For P2P Mode DIOs	22
9.3.	Processing a P2P Mode DIO	24
9.4.	Additional Processing of a P2P Mode DIO At An Intermediate Router	25
9.5.	Additional Processing of a P2P Mode DIO At The Target	26
9.6.	Processing a P2P-DR0 At An Intermediate Router	27
9.7.	Processing a P2P-DR0 At The Origin	29
10.	The P2P Discovery Reply Object Acknowledgement (P2P-DR0-ACK)	30
11.	Secure P2P-RPL Operation	31
12.	Packet Forwarding Along a Route Discovered Using P2P-RPL	32
13.	Interoperability with Core RPL	33
14.	Security Considerations	33
15.	IANA Considerations	35
15.1.	Additions to Mode of Operation	35
15.2.	Additions to RPL Control Message Options	35

15.3 . Additions to RPL Control Codes	36
16 . Acknowledgements	37
17 . References	37
17.1 . Normative References	37
17.2 . Informative References	38

Goyal, et al.

Expires September 21, 2013

[Page 2]

Internet-Draft

[draft-ietf-roll-p2p-rpl-17](#)

March 2013

Authors' Addresses	38
------------------------------	--------------------

1. Introduction

Targeting Low power and Lossy Networks (LLNs), the IPv6 Routing Protocol for LLNs (RPL) [[RFC6550](#)] provides paths along a Directed Acyclic Graph (DAG) rooted at a single router in the network. Establishment and maintenance of a DAG is performed by routers in the LLN using Destination-Oriented DAG (DODAG) Information Object (DIO) messages. When two arbitrary routers (neither of which is the DAG's root) need to communicate, the data packets are restricted to travel only along the links in the DAG. Such point-to-point (P2P) routing functionality may not be sufficient for several Home and Building Automation applications [[RFC5826](#)] [[RFC5867](#)] due to the following reasons:

- o The need to pre-establish routes: each potential destination in the network must declare itself as such ahead of the time a source needs to reach it.
- o The need to route only along the links in the DAG: A DAG is built to optimize the routing cost to reach the root. Restricting P2P routes to use only the in-DAG links may result in significantly suboptimal routes and severe traffic congestion near the DAG root.

This document describes an extension to core RPL (i.e., the RPL functionality described in [[RFC6550](#)]) that enables an IPv6 router in the LLN to discover routes to one or more IPv6 routers in the LLN "on demand". The discovered routes may not be the best available but are guaranteed to meet the specified routing metric constraints. Thus, such routes are considered "good enough" from the application's

perspective. This reactive P2P route discovery mechanism is henceforth referred to as P2P-RPL.

A mechanism to measure the end-to-end cost of an existing route is specified in [[I-D.ietf-roll-p2p-measurement](#)]. As discussed in [Section 4](#), measuring the end-to-end cost of an existing route may help decide whether to initiate the discovery of a better route using P2P-RPL and the metric constraints to be used for this purpose.

[1.1](#). Known Issues/Future Work

This document is presented as an Experimental specification to facilitate P2P-RPL's deployment in LLN scenarios where reactive P2P route discovery is considered useful or necessary. It is anticipated that, once sufficient operational experience has been gained, this specification will be revised to progress it on to the Standards Track. Experience reports regarding P2P-RPL implementation and deployment are encouraged particularly with respect to:

Goyal, et al.

Expires September 21, 2013

[Page 4]

Internet-Draft

[draft-ietf-roll-p2p-rpl-17](#)

March 2013

- o Secure P2P-RPL operation ([Section 11](#));
- o Rules governing Trickle operation ([Section 9.2](#));
- o Values in the default DODAG Configuration Option ([Section 6.1](#));
- o The RPLInstanceID reuse policy ([Section 6.1](#));
- o Utility and implementation complexity of allowing multiple Target addresses in a P2P-RPL route discovery.

[2](#). The Use Cases

One use case, common in home [[RFC5826](#)] and commercial building [[RFC5867](#)] environments, involves a device (say a remote control) that suddenly needs to communicate with another device (say a lamp) to which it does not already have a route (and whose network address it knows apriori). In this case, the remote control must be able to discover a route to the lamp "on demand".

Another use case, common in a 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. In this case, it is desirable to discover direct routes between various source-destination pairs that do not pass through the DAG's root.

Other use cases involve scenarios where energy or latency constraints are not satisfied by the P2P routes along an existing 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 [\[RFC6550\]](#) and [\[I-D.ietf-roll-terminology\]](#). This document introduces the following terms:

Origin : The IPv6 router initiating the P2P-RPL route discovery.

Target : The IPv6 router at the other end point of the P2P route(s)

to be discovered. A P2P-RPL route discovery can discover routes to multiple Targets at the same time.

Intermediate Router: An IPv6 router that is neither the Origin nor a Target.

Forward direction: The direction from the Origin to the Target.

Reverse direction: The direction from the Target to the Origin.

Forward Route: A route in the Forward direction.

Reverse Route: A route in the Reverse direction.

Bidirectional Route: A route that can be used in both Forward and

Reverse directions.

Ingress-only Interface: A network interface that can only receive packets.

Egress-only Interface: A network interface that can only send packets.

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.

RPL Security Configuration: The values for Counter is Time, Security Algorithm, Key Identifier Mode and Security Level fields, defined in [Section 6.1 of \[RFC6550\]](#), inside the Security section of a secure RPL control message.

[4.](#) Applicability

A route discovery using P2P-RPL may be performed by an Origin when no route exists between itself and the Target(s) or when the existing routes do not satisfy the application requirements. P2P-RPL is designed to discover Hop-by-hop or Source Routes to one or more Targets such that the discovered routes meet the specified constraints. In some application contexts, the constraints that the discovered routes must satisfy are intrinsically known or can be specified by the application. For example, an Origin that expects its Targets 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 the existing route to a Target to determine

the constraints. For example, an Origin that measures the total ETX along its current route to a Target to be 20 may use "ETX < x*20", where x is a fraction that the Origin decides, as the constraint. A mechanism to measure the cost of an existing route between two IPv6 routers is specified in [\[I-D.ietf-roll-p2p-measurement\]](#). If there is no existing route between the Origin and the Target(s) or the cost measurement for the existing routes fails, the Origin will have to guess the constraints to be 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.

P2P-RPL may result in discovery of better P2P routes than the ones available along a global 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 "distance" between the Origin and the Target (in terms of the routing metrics in use) and the prevalent conditions in the network. In general, a P2P-RPL route may be better than the one along a global DAG if the Origin and the Target are nearby. Similarly, a P2P-RPL route may not be much better than the one along a global DAG if the Origin and the Target are far apart. Note that, even when P2P-RPL routes are not much better than those along a global DAG, P2P-RPL routes may still be able to avoid congestion that might occur near the root if the routing takes place only along a global DAG. In general, the costs associated with a P2P-RPL route discovery (in terms of the control messages, mostly DIOs, generated) increases with the distance between the Origin and the Target. However, it is possible to limit the cost of route discovery by carefully setting the routing constraints, the Trickle parameters (that govern the DIO generation) and the time duration for which a router maintains its membership in the temporary DAG created for the route discovery. A network designer may take into consideration both the benefits (potentially better routes; no need to maintain routes proactively; avoid congestion near the global DAG's root) and costs when using P2P-RPL. The latency associated with a P2P-RPL route discovery again depends on the distance between the Origin and the Target and the Trickle parameters.

Like core RPL [[RFC6550](#)], P2P-RPL operation requires links to have bidirectional reachability. For this reason, the routers participating in a P2P-RPL route discovery must ensure that

- o Links that do not have bidirectional reachability do not become part of the route being discovered; and
- o IPv6 addresses belonging to Ingress-only (or Egress-only) Interfaces do not become part of the route being discovered.

This section contains a high level description of P2P-RPL.

A P2P-RPL route discovery takes place by forming a DAG rooted at the Origin. As is the case with core RPL, P2P-RPL uses IPv6 link-local multicast DIO messages to establish a DAG. However, unlike core RPL, this DAG is temporary in nature. The routes are discovered and installed while the DAG is alive. Once the specified duration of their membership in the DAG is over, the routers leave the DAG and hence the DAG ceases to exist. However, the installed routes are retained for their specified life time (which is different than the specified duration of a router's membership in the DAG) even though the DAG that caused their installation no longer exists. In P2P-RPL, the sole purpose of DAG creation is to discover routes to the Target(s) and DIOs serve as the route discovery messages. Each router joining the DAG determines a rank for itself in the DAG and ignores the subsequent DIOs received from lower (higher in numerical value) ranked neighbors. Thus, the route discovery messages propagate away from the Origin rather than return back to it. As in core RPL, DIO generation at a router is controlled by a Trickle timer [[RFC6206](#)] that allows a router to avoid generating unnecessary messages while providing protection against packet loss. P2P-RPL also uses the routing metrics [[RFC6551](#)], objective functions and packet forwarding framework [[RFC6554](#)][[RFC6553](#)] developed for core RPL.

An Origin may use P2P-RPL to discover routes to one or more Target(s) identified by one or more unicast/multicast addresses. P2P-RPL allows for the discovery of one Hop-by-hop Route or up to four Source Routes per Target. The discovered routes are guaranteed to meet the specified routing metric constraints but may not be the best available. P2P-RPL may fail to discover any route if the specified routing constraints are overly strict.

The Origin initiates a P2P-RPL route discovery by forming a temporary DAG rooted at itself. The DIOs used to create the temporary DAG are identified by a new Mode of Operation (P2P Route Discovery mode defined in [Section 6](#)). The DIOs listing the P2P Route Discovery mode as the Mode of Operation are henceforth referred to as the P2P mode DIOs. A P2P mode DIO always carries exactly one P2P Route Discovery Option (P2P-RDO, defined in [Section 7](#)) in which the Origin specifies the following information:

- o The IPv6 address of a Target. This could be a unicast address or a multicast address. Any additional Targets may be specified by including one or more RPL Target Options [[RFC6550](#)] inside the DIO.

- 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 per Target.
- o The desired number of routes (if Source Routes are being discovered).
- o Whether the Target(s) should send P2P Discovery Reply Object (P2P-DR0) messages (defined in [Section 8](#)) back to the Origin on receiving a DIO message. A P2P-DR0 message carries a discovered Source Route back to the Origin or establishes a Hop-by-hop Route between the Origin and the Target.

A P2P-RD0 also includes the best route from the Origin that the router, generating the P2P mode DIO, has seen so far.

A P2P mode DIO MAY also carry:

- o One or more Metric Container Options to specify:
 - * The relevant routing metrics.
 - * The constraints that the discovered route must satisfy. These constraints also limit how far the DIOs message may travel.
- o One or more RPL Target options to specify additional unicast or multicast Targets.

As the routers join the temporary DAG, they keep track of the best route(s) (so far from the Origin) they have seen and advertise these routes, along with the corresponding routing metrics, in their P2P mode DIOs. A router, including the Target(s), discards a received P2P mode 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 P2P mode DIO messages. A router may also discard a received P2P mode DIO if it does not wish to be a part of the discovered route due to limited resources or due to policy reasons.

When a Target receives a P2P mode DIO, it contains inside the P2P-RD0 a complete Source Route from the Origin to this Target. Since the links in the discovered route have bidirectional reachability ([Section 7](#)), the Target may use the discovered route to reach the Origin. Thus, a router that provides a particular service in the LLN (e.g. an outside temperature server) could initiate a P2P-RPL route discovery listing all its potential clients as Targets, thereby

allowing the clients to discover a Source Route back to the server. In this case, the Origin (the server) might want to disable the

generation of P2P-DR0 messages by the Targets (the clients). If the Origin has requested P2P-DR0 messages to be sent back, the Target may select the discovered route in the received DIO for further processing as described next. This document does not specify a particular method for the Target to use to select a route for further processing. Example 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 Route(s) are being discovered, the Target sends the selected Source Route(s) to the Origin via P2P-DR0 messages with one P2P-DR0 message carrying one discovered route. On receiving a P2P-DR0 message, the Origin stores the discovered route in its memory. This specification allows the Origin to discover up to four Source Routes per Target, thereby allowing the Origin to have sufficient ready-to-use alternatives should one or more of these routes fail. If a Hop-by-hop Route is being discovered, the Target sends a P2P-DR0 message containing the selected route to the Origin. The P2P-DR0 message travels back to the Origin along the selected route, establishing state for the Forward Route in the routers on the path.

The Target may request the Origin to acknowledge the receipt of a P2P-DR0 message by sending back a P2P-DR0 Acknowledgement (P2P-DR0-ACK) message (defined in [Section 10](#)). The Origin unicasts a P2P-DR0-ACK message to the Target. If the Target does not receive the requested P2P-DR0-ACK within a certain time interval of sending a P2P-DR0, it resends the P2P-DR0 message (up to a certain number of times) carrying the same route as before.

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" flag in the P2P-DR0 message to let the nodes in the LLN know about the completion of the route discovery process. The routers receiving such a P2P-DR0 should not generate any more DIOs for this temporary DAG. Neither should they process any received DIOs for this temporary DAG in future. However, such routers must still process

the P2P-DR0s received for this temporary DAG.

[6.](#) P2P Route Discovery Mode Of Operation

This section specifies a new RPL Mode of Operation (MOP), P2P Route Discovery Mode (or P2P mode, for short), with value TBD1. A DIO message, listing P2P mode as the MOP, is identified as performing a P2P-RPL route discovery by creating a temporary DAG. A P2P mode DIO

Goyal, et al.

Expires September 21, 2013

[Page 10]

Internet-Draft

[draft-ietf-roll-p2p-rpl-17](#)

March 2013

MUST carry exactly one P2P Route Discovery Option (P2P-RD0, specified in [Section 7](#)).

[6.1.](#) Setting a P2P Mode DIO

The Base Object in a P2P mode DIO message MUST be set in the following manner:

- o RPLInstanceID: RPLInstanceID MUST be a local value as described in [Section 5.1 of \[RFC6550\]](#). The Origin chooses the RPLInstanceID to be used for a particular route discovery in accordance with the following rules:
 - * The Origin SHOULD NOT reuse an RPLInstanceID for a route discovery if some routers might still maintain membership in the DAG the Origin had initiated for the previous route discovery using this RPLInstanceID. As described in [Section 7](#), a router's membership in a DAG created for a P2P-RPL route discovery lasts for the time duration (say 'l' seconds) indicated by the L field inside the P2P-RD0. In general, there is no upper bound on the time duration by when all the routers have left the DAG created for a P2P-RPL route discovery. In the specific case where the discovered route must be at most 'n' hops in length, all the routers must have left the DAG "(n+1)*l" seconds after its initiation by the Origin. In practice, all the routers should have joined the DAG within 'l' seconds of its initiation (since the route discovery must complete while the Origin still belongs to the DAG) and hence all the routers should have left the DAG within "2*l" seconds of its initiation. Hence, it is usually sufficient that the Origin wait for twice the duration indicated by the L field inside the P2P-RD0 used for the previous route discovery before

reusing the RPLInstanceID for a new route discovery. Individual P2P-RPL deployments are encouraged to share their experience with various RPLInstanceID reuse policies to help guide the development of standards track version of the protocol.

- * When initiating a new route discovery to a particular Target, the Origin MUST NOT reuse the RPLInstanceID used in a previous route discovery to this Target if the state created during the previous route discovery might still exist in some routers. Note that it is possible that the previous route discovery did not succeed yet some routers still ended up creating state. The Default Lifetime and Lifetime Unit parameters in the DODAG Configuration Option specify the lifetime of the state the routers, including the Origin and the Target, maintain for a Hop-by-hop or a Source Route discovered using P2P-RPL. Suppose

this lifetime is 'X' seconds. As discussed above, any state created during the previous route discovery was likely created within "2*l" seconds of its initiation. Hence, it is sufficient that the Origin lets a time duration equal to "X+2*l" seconds pass since the initiation of the previous route discovery before initiating a new route discovery to the same Target using the same RPLInstanceID.

- o Version Number: MUST be set to zero. The temporary DAG used for P2P-RPL route discovery does not exist long enough to have new versions.
- o Grounded (G) Flag: This flag MUST be set to one. Unlike a global RPL instance, the concept of a floating DAG, used to provide connectivity within a sub-DAG detached from a grounded DAG, does not apply to a local RPL instance. Hence, an Origin MUST always set the G flag to one when initiating a P2P-RPL route discovery. Further, clause 3 of [Section 8.2.2.2 in \[RFC6550\]](#) does not apply and a node MUST NOT initiate a new DAG if it does not have any parent left in a P2P-RPL DAG.
- o Mode of Operation (MOP): MUST be set to TBD1, corresponding to P2P Route Discovery mode.
- o DTSN: MUST be set to zero on transmission and ignored on

reception.

- o DODAGPreference (Prf): This field MUST be set to zero (least preferred).
- o DODAGID: This field MUST be set to an IPv6 address of the Origin.
- o The other fields in the DIO Base Object can be set in the desired fashion as per the rules described in [[RFC6550](#)].

A received P2P mode DIO MUST be discarded if it does not follow the above-listed rules regarding the RPLInstanceID, Version Number, G flag, MOP and Prf fields inside the base object.

The DODAG Configuration Option, inside a P2P mode DIO MUST be set in the following manner:

- o The Origin MUST set the MaxRankIncrease parameter to zero to disable local repair of the temporary DAG. A received P2P mode DIO MUST be discarded if the MaxRankIncrease parameter inside the DODAG Configuration Option is not zero.

- o The Origin SHOULD set the Trickle parameters (DIOIntervalDoublings, DIOIntervalMin, DIORedundancyConstant) as recommended in [Section 9.2](#).
- o The Origin sets the Default Lifetime and Lifetime Unit parameters to indicate the lifetime of the state the routers, including the Origin and the Target(s), maintain for a Hop-by-hop or a Source Route discovered using P2P-RPL.
- o The Origin sets the other fields in the DODAG Configuration Option, including the OCP identifying the Objective function, in the desired fashion as per the rules described in [[RFC6550](#)].
- o An Intermediate Router (or a Target) MUST set various fields in the DODAG Configuration Option in the outgoing P2P mode DIOs to the values they had in the incoming P2P mode DIOs for this DAG.

A default DODAG Configuration Option comes in effect if a P2P mode

DIO does not carry an explicit one. The default DODAG Configuration Option has the following parameter values:

- o Authentication Enabled: 0
- o DIOIntervalMin: 6, which translates to 64ms as the value for Imin parameter in Trickle operation. This value is roughly one order of magnitude larger than the typical transmission delay on IEEE 802.15.4 links and corresponds to the recommendation in [Section 9.2](#) for well-connected topologies.
- o DIORedundancyConstant: 1. See the discussion in [Section 9.2](#).
- o MaxRankIncrease: 0 (to disable local repair of the temporary DAG).
- o Default Lifetime: 0xFF, to correspond to infinity.
- o Lifetime Unit: 0xFFFF, to correspond to infinity.
- o Objective Code Point: 0, i.e., 0F0 [[RFC6552](#)] is the default objective function.
- o The remaining parameters have default values as specified in [[RFC6550](#)].

Individual P2P-RPL deployments are encouraged to share their experience with these default values to help guide the development of standards track version of the protocol.

The routing metrics and constraints [[RFC6551](#)] used in P2P-RPL route

discovery are included in one or more Metric Container Options [[RFC6550](#)] inside the P2P mode DIO. Note that a DIO need not include a Metric Container if 0F0 is the objective function in effect. In that case, a P2P mode DIO may still specify an upper limit on the maximum rank, that a router may have in the temporary DAG, inside the P2P-RD0.

A P2P mode DIO:

- o MUST carry one (and only one) P2P-RD0. The P2P-RD0 allows for the specification of one unicast or multicast address for the Target.

A received P2P mode DIO MUST be discarded if it does not contain exactly one P2P-RD0.

- o MAY carry one or more RPL Target Options to specify additional unicast/multicast addresses for the Target. If a unicast address is specified, it MUST be a global address or a unique local address.
- o MAY carry one or more Metric Container Options to specify routing metrics and constraints.
- o MAY carry one or more Route Information Options [[RFC6550](#)]. In the context of P2P-RPL, a Route Information Option advertizes to the Target(s) the Origin's connectivity to the prefix specified in the option.

An RPL Option, besides the ones listed above, **MUST** be ignored when found inside a received P2P mode DIO and **MUST NOT** be included in the P2P mode DIOs the receiving router generates.

In accordance with core RPL, a P2P mode DIO MUST propagate via link-local multicast. The IPv6 source address in a P2P mode DIO MUST be a link-local address and the IPv6 destination address MUST be the link-local multicast address all-RPL-nodes [[RFC6550](#)]. A P2P mode DIO MUST be transmitted on all interfaces the router has in this RPL domain [[I-D.ietf-roll-terminology](#)].

7. P2P Route Discovery Option (P2P-RD0)

This section defines a new RPL control message option: the P2P Route Discovery Option (P2P-RDO).

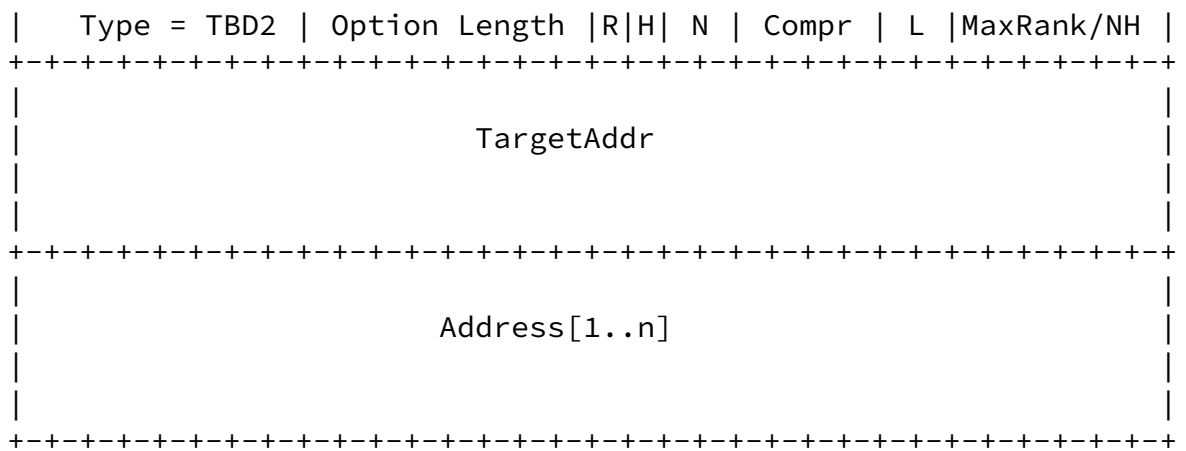


Figure 1: Format of P2P Route Discovery Option (P2P-RD0)

The format of a P2P Route Discovery Option (P2P-RD0) is illustrated in Figure 1. A P2P mode DIO and a P2P-DR0 (defined in [Section 8](#)) message MUST carry exactly one P2P-RD0. A P2P-RD0 consists of the following fields:

- o Option Type: TBD2.
- o Option Length: 8-bit unsigned integer, representing the length in octets of the option, not including the Option Type and Option Length fields.
- o Reply (R): The Origin sets this flag to one to allow the Target(s) to send P2P-DR0 messages back to the Origin. If this flag is zero, a Target MUST NOT generate any P2P-DR0 message.
- o Hop-by-hop (H): This flag is valid only if the R flag is set to one. The Origin sets this flag to one if it desires Hop-by-hop Routes. The Origin sets this flag to zero if it desires Source Routes. This specification allows for the establishment of one Hop-by-hop route or up to four Source Routes per Target. The Hop-by-hop Route is established in the Forward direction, i.e. from the Origin to the Target. This specification does not allow for the establishment of Hop-by-hop Routes in the Reverse direction.
- o Number of Routes (N): This field is valid only if the R flag is one and H flag is zero, i.e. the Targets are allowed to generate P2P-DR0 messages carrying discovered Source Routes back to the Origin. In this case, the value in the N field plus one indicates

the number of Source Routes that each Target should convey to the Origin. When Hop-by-hop Routes are being discovered, the N 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 zero 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 exact duration a router joining the temporary DAG, including the Origin and the Target(s), MUST maintain its membership in the DAG. A router MUST leave the temporary DAG once the time elapsed since it joined reaches the value indicated by this field. The mapping between the value in this field and the duration of the router's membership in the temporary DAG is as follows:

- * 0x00: 1 second;
- * 0x01: 4 seconds;
- * 0x02: 16 seconds;
- * 0x03: 64 seconds;

The Origin sets this field based on its expectation regarding the time required for the route discovery to complete, which includes the time required for the DIOs to reach the Target(s) and the P2P-DROs to travel back to the Origin. The time required for the DIOs to reach the Target(s) would in turn depend on the Trickle parameters (Imin and the redundancy constant) as well as the expected distance (in terms of hops and/or ETX) to the Target(s). While deciding the value in this field, the Origin should also take in account the fact that all routers joining the temporary DAG would need to stay in the DAG for this much time.

- o MaxRank/NH:
 - * When a P2P-RDO is included in a P2P mode DIO, this field indicates the upper limit on the integer portion of the rank (calculated using the DAGRank() macro defined in [[RFC6550](#)]) that a router may have in the temporary DAG being created. An Intermediate Router MUST NOT join a temporary DAG being created by a P2P mode DIO if the integer portion of its rank would be equal to or higher (in numerical value) than the MaxRank limit. A Target can join the temporary DAG at a rank whose integer portion is equal to the MaxRank. A router MUST discard a

received P2P mode DIO if the integer part of the advertized

rank equals or exceeds the MaxRank limit. A value 0 in this field indicates that the MaxRank is infinity.

- * When a P2P-RD0 is included in a P2P-DR0 message, this field indicates the index of the next hop address inside the Address vector.
- o TargetAddr: An IPv6 address of the Target after eliding Compr number of prefix octets. When the P2P-RD0 is included in a P2P mode DIO, this field may contain a unicast address or a multicast address. If a unicast address is specified, it MUST be a global address or a unique local address. Any additional Target addresses can be specified by including one or more RPL Target Options [[RFC6550](#)] in the DIO. When the P2P-RD0 is included in a P2P-DR0, this field MUST contain a unicast global or unique local IPv6 address of the Target generating the P2P-DR0.
- o Address[1..n]: A vector of IPv6 addresses representing a complete route so far in the Forward direction:
 - * Each element in the Address vector has size $(16 - \text{Compr})$ octets and MUST contain a valid global or unique local IPv6 address with first Compr octets elided.
 - * The total number of elements inside the Address vector is given by $n = (\text{Option Length} - 2 - (16 - \text{Compr})) / (16 - \text{Compr})$.
 - * The IPv6 address that a router adds to the vector MUST belong to the interface on which the router received the DIO containing this P2P-RD0. Further, this interface MUST NOT be an Ingress-only Interface. This allows the route accumulated in the Address vector to be a Bidirectional Route that can be used by a Target to send a P2P-DR0 message 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 any of its addresses do not already exist in the vector. A Target specifying a complete route in the Address vector MUST ensure that the vector does not contain any address more than once.

Goyal, et al.

Expires September 21, 2013

[Page 17]

Internet-Draft

[draft-ietf-roll-p2p-rpl-17](#)

March 2013

- * The Address vector MUST NOT contain any multicast addresses.

8. The P2P Discovery Reply Object (P2P-DR0)

This section defines two new RPL Control Message types, the P2P Discovery Reply Object (P2P-DR0), with code TBD3, and the Secure P2P-DR0, with code TBD4. A P2P-DR0 serves one of the following functions:

- o Carry a discovered Source Route from a Target to the Origin;
- o Establish a Hop-by-hop Route as it travels from a Target to the Origin.

A P2P-DR0 message can also serve the function of letting the routers in the LLN know that a P2P-RPL route discovery is complete and no more DIO messages need to be generated for the corresponding temporary DAG. A P2P-DR0 message MUST carry one (and only one) P2P-RD0 whose TargetAddr field MUST contain a unicast IPv6 address of the Target that generates the P2P-DR0. A P2P-DR0 message MUST travel from the Target to the Origin via link-local multicast along the route specified inside the Address vector in the P2P-RD0, as included in the P2P-DR0. The IPv6 source address in a P2P-DR0 message MUST be a link-local address and the IPv6 destination address MUST be the link-local multicast address all-RPL-nodes [[RFC6550](#)]. A P2P-DR0 message MUST be transmitted on all interfaces the router has in this RPL domain [[I-D.ietf-roll-terminology](#)].

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	2	3	4	5	6	7	8	9
RPLInstanceID										Version										S	A	Seq	Reserved																

and DODAGID fields) had the Stop flag set to one. When set to zero, this flag does not imply any thing and MUST be ignored on reception.

- o Ack Required (A): This flag, when set to one by the Target, indicates that the Origin MUST unicast a P2P-DR0-ACK message (defined in [Section 10](#)) to the Target when it receives the P2P-DR0.
- o Sequence Number (Seq): This 2-bit field indicates the sequence number for the P2P-DR0. This field is relevant when the A flag is set to one, i.e., the Target requests an acknowledgement from the Origin for a received P2P-DR0. The Origin includes the RPLInstanceID, the DODAGID and the Sequence Number of the received P2P-DR0 inside the P2P-DR0-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 P2P-DR0 message:
 - * MUST carry one (and only one) P2P-RD0 that MUST specify a complete route between the Target and the Origin. A received P2P-DR0 message MUST be discarded if it does not contain exactly one P2P-RD0.
 - * MAY carry one or more Metric Container Options that contains the aggregated routing metrics values for the route specified in P2P-RD0.

An RPL Option, besides the ones listed above, MUST be ignored when found inside a received P2P-DR0 message.

[8.1](#). Secure P2P-DR0

A Secure P2P-DR0 message follows the format in Figure 7 of [[RFC6550](#)], where the base format is the base P2P-DR0 shown in Figure 2.

[8.2.](#) Setting a P2P-RD0 Carried in a P2P Discovery Reply Object

A P2P Discovery Reply Object MUST carry one (and only one) P2P-RD0, which MUST be set as defined in [Section 7](#). Specifically, the following fields MUST be set as specified next:

- o Reply (R): This flag MUST be set to zero on transmission and ignored on reception.
- o Hop-by-Hop (H): The H flag in the P2P-RD0 included in a P2P-DR0 message MUST have the same value as the H flag in the P2P-RD0 inside the corresponding DIO message.
- o Number of Routes (N): This field MUST be set to zero on transmission and ignored on reception.
- o Life Time (L): This field MUST be set to zero on transmission and ignored on reception.
- o MaxRank/NH: This field indicates the index of the next hop address in the Address vector. When a Target generates a P2P-DR0 message, the NH field is set to $n = (\text{Option Length} - 2 - (16 - \text{Compr})) / (16 - \text{Compr})$.

- o TargetAddr: This field MUST contain a unicast global or unique local IPv6 address of the Target generating the P2P-DR0.
- o Address[1..n]: 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.

[9.](#) P2P-RPL Route Discovery By Creating a Temporary DAG

This section details the P2P-RPL route discovery operation.

[9.1.](#) Joining a Temporary DAG

All the routers participating in a P2P-RPL route discovery, including the Origin and the Target(s), MUST join the temporary DAG being created for the purpose. When a router joins a temporary DAG advertized by a P2P mode DIO, it MUST maintain its membership in the temporary DAG for the duration indicated by the L field inside the P2P-RD0. The only purpose of a temporary DAG's existence is to facilitate the P2P-RPL route discovery process. The temporary DAG MUST NOT be used to route data packets. In other words, joining a temporary DAG does not allow a router to provision routing table entries listing the router's parents in the temporary DAG as the next hops (i.e., the last bullet point in [Section 3.2.8 of \[RFC6550\]](#) is not applicable when the DAG is a temporary DAG created for the purpose of a P2P-RPL route discovery.).

Given the nature of a temporary DAG created for a P2P-RPL route discovery, this document disallows the solicitation of P2P mode DIOs using DODAG Information Solicitation (DIS) messages as described in [\[RFC6550\]](#). A router participating in a P2P-RPL route discovery MUST NOT reset its Trickle timer that controls the transmission of P2P mode DIOs in response to a multicast DIS. Also, the router MUST NOT send a P2P mode DIO in response to a unicast DIS. In other words, the rules in [Section 8.3 of \[RFC6550\]](#) regarding a router's response to a multicast/unicast DIS are not applicable for P2P mode DIOs.

A router MUST detach from the temporary DAG created for a P2P-RPL route discovery once the duration of its membership in the DAG has reached the value indicated by the L field inside the P2P-RD0. After receiving a P2P-DR0 with the Stop flag set to one, a router SHOULD NOT send or process any more DIOs for this temporary DAG and SHOULD also cancel any pending DIO transmission.

[9.2.](#) Trickle Operation For P2P Mode DIOs

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 LLNs. When controlling the transmissions of a P2P mode DIO, a Trickle timer SHOULD follow the following rules:

- o The receipt of a P2P mode DIO, that allows the router to advertise a better route (in terms of the routing metrics and the OF in use) than before, is considered "inconsistent" and hence resets the Trickle timer. Note that the first receipt of a P2P mode DIO advertising a particular temporary DAG is always considered an "inconsistent" event.
- o The receipt of a P2P mode DIO from a parent in the temporary DAG is considered neither "consistent" nor "inconsistent" if it does not allow the router to advertise a better route than before. Thus, the receipt of such DIOs has no impact on the Trickle operation. Note that this document does not impose any requirements on how a router might choose its parents in the temporary DAG.
- o The receipt of a P2P mode DIO is considered "consistent" if the source of the DIO is not a parent in the temporary DAG and either of the following conditions is true:
 - * The DIO advertises a better route than the router but does not allow the router to advertise a better route itself; or
 - * The DIO advertises a route as good as the route (to be) advertised by the router.

Note that the Trickle algorithm's DIO suppression rules are in effect at all times. Hence, a P2P-RPL router may suppress a DIO transmission even if it has not made any DIO transmission yet.

- o The receipt of a P2P mode DIO, that advertises a worse route than what the router advertises (or would advertise when it gets a chance to generate its DIO), is considered neither "consistent" nor "inconsistent", i.e., the receipt of such a DIO has no impact on the Trickle operation.
- o The Imin parameter SHOULD be set taking in account the connectivity within the network. For highly connected networks, a small Imin value (of the order of the typical transmission delay

for a DIO) may lead to congestion in the network as a large number of routers reset their Trickle timers in response to the first receipt of a DIO from the Origin. These routers would generate their DIOs within Imin interval and cause additional routers to reset their trickle timers and generate more DIOs. Thus, for highly connected networks, the Imin parameter SHOULD be set to a value at least one order of magnitude larger than the typical transmission delay for a DIO. For sparsely connected networks, the Imin parameter can be set to a value that is a small multiple of the typical transmission delay for a DIO. Note that the Imin value has a direct impact on the time required for a P2P-RPL route discovery to complete. In general, the time required for a P2P-RPL route discovery would increase approximately linearly with the value of the Imin parameter. Since the route discovery must complete while the Origin still belongs to the temporary DAG created for the purpose, the Origin should set the time duration a router maintains its membership in the temporary DAG (indicated by the L field inside the P2P-RD0) to a large enough value taking in account the Imin value as well as the expected distance (in terms of hops and/or ETX) to the Target(s).

- o The Imax parameter SHOULD be set to a large value (several orders of magnitude higher than the Imin value) and is unlikely to be critical for P2P-RPL operation. This is because the first receipt of a P2P mode DIO for a particular temporary DAG is considered an inconsistent event and would lead to resetting of Trickle timer duration to the Imin value. Given the temporary nature of the DAGs used in P2P-RPL, Trickle timer may not get a chance to increase much.
- 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. This setting for the redundancy constant is designed to reduce the number of messages generated during a route discovery process and is suitable for the environments with low or moderate packet loss rates. However, this setting may result in an increase in the time required for the route discovery process to complete. A higher value for the redundancy constant may be more suitable in
 - * Environments with high packet loss rates; or
 - * Deployments where the time required for the route discovery process to complete needs to be as small as possible; or
 - * Deployments where specific destinations are reachable only through specific intermediate routers (and hence these intermediate routers should not suppress their DIOs).

A particular deployment should take in account the above mentioned factors when deciding the value of the redundancy constant.

Individual P2P-RPL deployments are encouraged to share their experience with these rules to help guide the development of standards track version of the protocol. Applicability Statements that specify the use of P2P-RPL MUST provide guidance for setting Trickle parameters, particularly Imin and the redundancy constant.

[9.3.](#) Processing a P2P Mode DIO

The rules for DIO processing and transmission, described in [Section 8](#) of RPL [[RFC6550](#)], apply to P2P mode DIOs as well except as modified in this document. In particular, in accordance with [Section 8.2.3](#) of RPL [[RFC6550](#)], a received P2P mode DIO MUST be discarded if it is malformed according to the rules specified in this document and in [[RFC6550](#)].

The following rules for processing a received P2P mode DIO apply to both Intermediate Routers and the Target.

A router SHOULD discard a received P2P mode DIO with no further processing if it does not have bidirectional reachability with the neighbor that generated the received DIO. 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 calculate the values of the link-level routing metrics included in the received DIO taking in account the metric's value in both Forward and Reverse directions. Bidirectional reachability along a discovered route allows the Target to use this route to reach the Origin. In particular, the P2P-DRO messages travel from the Target to the Origin along a discovered route.

A router MUST discard a received P2P mode DIO with no further processing:

- o If the DIO advertises INFINITE_RANK as defined in [Section 17 of \[\[RFC6550\]\(#\)\]](#).
- o If the integer part of the rank advertised in the DIO equals or exceeds the MaxRank limit listed in the P2P Route Discovery Option.

- o If the routing metric values do not satisfy one or more of the mandatory route constraints listed in the DIO or if the router cannot evaluate the mandatory route constraints, e.g., if the router does not support the metrics used in the constraints.

- o If the router previously received a P2P-DR0 message with the same RPLInstanceID and DODAGID as the received DIO and with the Stop flag set to one.

The router MUST check the Target addresses listed in the P2P-RD0 and any RPL Target Options included in the received DIO. If one of its IPv6 addresses is listed as a Target address or if it belongs to the multicast group specified as one of the Target addresses, the router considers itself a Target and processes the received DIO as specified in [Section 9.5](#). Otherwise, the router considers itself an Intermediate Router and processes the received DIO as specified in [Section 9.4](#).

[9.4](#). Additional Processing of a P2P Mode DIO At An Intermediate Router

An Intermediate Router MUST discard a received P2P mode DIO with no further processing

- o if the DIO is received on an Ingress-only Interface; or
- o if the receiving interface does not have a global or unique local IPv6 address configured with the address prefix implied by the Compr field in the P2P-RD0 inside the received DIO; or
- o if the router can not uniquely identify the address prefix implied by the Compr field in the P2P-RD0 (this might happen if the receiving interface has multiple global/unique-local IPv6 addresses, each configured with a different address prefix); or
- o if adding its IPv6 address to the route in the Address vector inside the P2P-RD0 would result in the route containing multiple addresses belonging to this router.

On receiving a P2P mode DIO, an Intermediate Router MUST do the following. The router 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 9.2](#). Note that the route comparison in a P2P-RPL route discovery is performed using the parent selection rules of the OF in use as specified in [Section 14](#) of RPL [[RFC6550](#)]. If the received DIO would allow the router to advertise a better route, the router MUST add a unicast IPv6 address of the receiving interface (after eliding Compr prefix octets) to the route in the Address vector inside the P2P-RDO and remember this route for inclusion in its future DIOs.

When an Intermediate Router adds an IPv6 address to a route, it MUST

ensure that

- o the IPv6 address is a unicast global or unique local IPv6 address assigned to the interface on which the DIO containing the route was received;
- o the IPv6 address was configured with the address prefix implied by the Compr field in the P2P-RDO inside the received DIO;

To improve the diversity of the routes being discovered, an Intermediate Router SHOULD keep track of multiple routes (as long as all these routes are the best seen so far), one of which SHOULD be selected in a uniform random manner for inclusion in the P2P-RDO inside the router's next DIO. Note that the route accumulation in a P2P mode DIO MUST take place even if the Origin does not want any P2P-RDO messages to be generated (i.e., the R flag inside the P2P-RDO is set to zero). This is because the Target may still be able to use the accumulated route as a source route to reach the Origin.

[9.5](#). Additional Processing of a P2P Mode DIO At The Target

The Target MAY remember the discovered route contained in the P2P-RDO in the received DIO for use as a Source Route to reach the Origin. The lifetime of this Source Route is specified by the Default Lifetime and Lifetime Unit parameters inside the DODAG Configuration Option currently in effect. This lifetime can be extended (or shortened) appropriately following a hint from an upper-layer protocol.

If the Reply flag inside the P2P-RD0 in the received DIO is set to one, the Target MUST select one or more discovered routes and send one or more P2P-DR0 messages, carrying one discovered route each, back to the Origin. If the H flag inside the P2P-RD0 is set to one, the Target needs to select one route and send a P2P-DR0 message along this route back to the Origin. As this P2P-DR0 message travels back to the Origin, the routers on the path establish hop-by-hop routing state, thereby establishing a Hop-by-hop Route in the Forward direction. If the H flag is set to zero, the number of Source Routes to be selected (and the number of P2P-DR0 messages to be sent back) is given by one plus the value of the N field in the P2P-RD0. The Target may select the discovered route inside the received DIO as (one of) the route(s) that would be carried inside a P2P-DR0 message back to the Origin. This document does not prescribe a particular method for the Target to select the routes. Example methods include selecting each route that meets the specified routing constraints until the desired number have been selected or selecting the best routes discovered over a certain time period. If multiple routes are to be selected, the Target SHOULD avoid selecting routes that have

large segments in common.

If the Target selects the route contained in the P2P-RD0 in the received DIO, it sends a P2P-DR0 message back to the Origin (identified by the DODAGID field in the DIO). The P2P-DR0 message MUST include a P2P-RD0 that contains the selected route inside the Address vector. Various fields inside the P2P-RD0 MUST be set as specified in [Section 8.2](#). The Target MAY set the A flag inside the P2P-DR0 message to one if it desires the Origin to send back a P2P-DR0-ACK message on receiving the P2P-DR0. In this case, the Target waits for P2P_DR0_ACK_WAIT_TIME duration for the P2P-DR0-ACK message to arrive. Failure to receive the P2P-DR0-ACK message within this time duration causes the Target to retransmit the P2P-DR0 message. The Target MAY retransmit the P2P-DR0 message in this fashion up to MAX_P2P_DR0_RETRANSMISSIONS times. Both P2P_DR0_ACK_WAIT_TIME and MAX_P2P_DR0_RETRANSMISSIONS are configurable parameters to be decided based on the characteristics of individual deployments. Note that all P2P-DR0 transmissions and retransmissions MUST take place while the Target is still a part of the temporary DAG created for the route discovery. A Target MUST NOT transmit a P2P-DR0 if it no longer belongs to this DAG.

The Target MAY set the Stop flag inside the P2P-DR0 message to one if

- o this router is the only Target specified in the corresponding DIO, i.e., the corresponding DIO specified a unicast address of the router as the TargetAddr inside the P2P-RD0 with no additional Targets specified via RPL Target Options; and
- o the Target has already selected the desired number of routes.

The Target MAY include a Metric Container Option in the P2P-DR0 message. This Metric Container contains the end-to-end routing metric values for the route specified in the P2P-RD0. The Target MUST transmit the P2P-DR0 message via a link-local multicast.

A Target MUST NOT forward a P2P mode DIO any further if no other Targets are to be discovered, i.e., if a unicast IPv6 address (of this Target) is specified as the TargetAddr inside the P2P-RD0 and no additional Targets are specified via RPL Target Options inside the DIOs for this route discovery. Otherwise, the Target MUST generate DIOs for this route discovery as an Intermediate Router would.

[9.6.](#) Processing a P2P-DR0 At An Intermediate Router

If the DODAGID field in the received P2P-DR0 does not list a router's own IPv6 address, the router considers itself an Intermediate Router and MUST process the received message in the following manner:

- o The router MUST discard the received P2P-DR0 with no further processing if it does not belong to the temporary DAG identified by the RPLInstanceID and the DODAGID fields in the P2P-DR0.
- o If the Stop flag inside the received P2P-DR0 is set to one, the router SHOULD NOT send or receive any more DIOs for this temporary DAG and SHOULD cancel any pending DIO transmission.
- o The router MUST ignore any Metric Container Options contained in the P2P-DR0 message.
- o If Address[NH] element inside the P2P-RD0 lists the router's own unicast IPv6 address, the router is a part of the route carried in the P2P-RD0. In this case, the router MUST do the following:

- * To prevent loops, the router MUST discard the P2P-DR0 message with no further processing if the Address vector in the P2P-RD0 includes multiple IPv6 addresses assigned to the router's interfaces.
- * If the H flag inside the P2P-RD0 is one, the router MUST store the state for the Forward Hop-by-hop route carried inside the P2P-RD0. This state consists of:
 - + The RPLInstanceID and the DODAGID fields of the P2P-DR0.
 - + The route's destination, the Target (identified by TargetAddr field inside P2P-RD0).
 - + The IPv6 address of the next hop, Address[NH+1] (unless NH value equals the number of elements in the Address vector, in which case the Target itself is the next hop).

This Hop-by-hop routing state MUST expire at the end of the lifetime specified by the Default Lifetime and Lifetime Unit parameters inside the DODAG Configuration Option used in P2P mode DIOs for this route discovery.

- * 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 P2P-DR0 and the next hop information in the state does not match the next hop indicated in the received P2P-DR0, the router MUST discard the P2P-DR0 message with no further processing. Note that this situation would occur in the following two cases:
 - + When the route listed in the Address vector inside the P2P-RD0 contains a previously undetected loop. In this case,

the rule above causes the P2P-DR0 messages to be discarded.

- + When a Hop-by-hop Route between the Origin and the Target, previously established using the same RPLInstanceID and DODAGID as the route currently being established, still exists and at least partially overlaps the route currently being established.

- * The router MUST decrement the NH field inside the P2P-RD0 and send the P2P-DRO message further via link-local multicast.

[9.7.](#) Processing a P2P-DRO At The Origin

When a router receives a P2P-DRO message that lists its IPv6 address in the DODAGID field, the router recognizes itself as the Origin for the corresponding P2P-RPL route discovery, notes the Target that originated this message (from the TargetAddr field inside the P2P-RD0) and processes the message in the following manner:

- o The Origin MUST discard the received P2P-DRO with no further processing if it no longer belongs to the temporary DAG identified by the RPLInstanceID and the DODAGID fields in the P2P-DRO.
- o If the Stop flag inside the received P2P-DRO is set to one, the Origin SHOULD NOT generate any more DIOs for this temporary DAG and SHOULD cancel any pending DIO transmission.
- o If the P2P-RD0 inside the P2P-DRO has the H flag set to 0, the Address vector inside the P2P-RD0 contains a Source Route to this Target. The Origin MUST set the lifetime of this Source Route to the value specified by the Default Lifetime and Lifetime Unit parameters inside the DODAG Configuration Option in the P2P mode DIOs used for this route discovery. This lifetime could be extended (or shortened) appropriately following a hint from an upper-layer protocol.
- o If the P2P-RD0 inside the P2P-DRO has the H flag set to 1, the P2P-DRO message is establishing a Hop-by-hop Route to this Target and the Origin MUST store in its memory the state for this Hop-by-hop Route in the manner described in [Section 9.6](#). This Hop-by-hop routing state MUST expire at the end of the lifetime specified by the Default Lifetime and Lifetime Unit parameters inside the DODAG Configuration Option used in P2P mode DIOs for this route discovery. The standards track version of P2P-RPL may consider specifying a signaling mechanism that will allow the Origin to extend (or shorten) the lifetime of a P2P-RPL Hop-by-hop Route following a suitable hint from an upper-layer protocol.

- o If the received P2P-DR0 message contains one or more Metric Container Options, 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-RPL route discovery to this Target.
- o If the A flag is set to one in the received P2P-DR0 message, the Origin MUST generate a P2P-DR0-ACK message as described in [Section 10](#) and unicast the message to the Target. The Origin MAY use the route just discovered to send the P2P-DR0-ACK message to the Target. [Section 12](#) describes how a packet may be forwarded along a Source/Hop-by-hop Route discovered using P2P-RPL.

[10.](#) The P2P Discovery Reply Object Acknowledgement (P2P-DR0-ACK)




Figure 3: Format of the base P2P Discovery Reply Object Acknowledgement (P2P-DR0-ACK)

A P2P-DR0 message may fail to reach the Origin due to a number of reasons. Unlike the DIO messages that benefit from Trickle-controlled retransmissions, the P2P-DR0 messages are prone to loss due to unreliable packet transmission in LLNs. Since a P2P-DR0 message travels via link-local multicast, it cannot use link-level acknowledgements to improve the reliability of its transmission. Also, an Intermediate Router may drop the P2P-DR0 message (e.g., because of its inability to store the state for the Hop-by-hop Route the P2P-DR0 is establishing). To protect against the potential failure of a P2P-DR0 message to reach the Origin, the Target MAY request the Origin to send back a P2P-DR0 Acknowledgement (P2P-DR0-ACK) message on receiving a P2P-DR0 message. Failure to receive such an acknowledgement within the P2P_DR0_ACK_WAIT_TIME interval of sending the P2P-DR0 message forces the Target to resend the message (as described in [Section 9.5](#)).

This section defines two new RPL Control Message types: P2P-DR0

Acknowledgement (P2P-DR0-ACK; with code TBD5) and Secure P2P-DR0-ACK (with code TBD6). A P2P-DR0-ACK message MUST travel as a unicast message from the Origin to the Target. The IPv6 source and destination addresses used in a P2P-DR0-ACK message MUST be global or unique local. The format of a base P2P-DR0-ACK message is shown in Figure 3. Various fields in a P2P-DR0-ACK message MUST have the same values as the corresponding fields in the P2P-DR0 message. The field marked as "Reserved" MUST be set to zero on transmission and MUST be ignored on reception. A Secure P2P-DR0-ACK message follows the format in Figure 7 of [RFC6550], where the base format is same as the base P2P-DR0-ACK shown in Figure 3.

11. Secure P2P-RPL Operation

Each RPL control message type, including the ones defined in this document, has a secure version. A secure RPL control message is identified by the value 1 in the most significant bit of the Code field. Each secure RPL control message contains a security section (see Figure 7 of [RFC6550]), whose contents are described in [Section 6.1 of \[RFC6550\]](#). Sections [6.1](#), [10](#) and [19](#) of [RFC6550] describe core RPL's security apparatus. These sections are applicable to P2P-RPL's secure operation as well except as constrained in this section.

Core RPL allows a router to decide locally on a per-packet basis whether to use security and if yes what Security Configuration (see definition in [Section 3](#)) to use (the only exception being the requirement to send a Secure DIO in response to a Secure DIS; see [Section 10.2 of \[RFC6550\]](#)). In contrast, this document requires routers participating in a P2P-RPL route discovery to follow the Origin's lead regarding security. The Origin decides whether to use security and the particular Security Configuration to be used for this purpose. All the routers participating in this route discovery MUST generate only secure control messages if the Origin decides so and MUST use for this purpose the Security Configuration that the Origin chose. The Origin MUST NOT set the "Key Identifier Mode" field inside the chosen Security Configuration to value 1 since this setting indicates the use of a per-pair key which is not suitable for securing messages that travel by (link local) multicast (e.g. DIOs) or that travel over multiple hops (e.g. P2P-DR0s). The Origin MUST use the chosen Security Configuration to secure all the control messages (DIOs and P2P-DR0-ACKs) it generates.

A router MUST NOT join the temporary DAG being created for a P2P-RPL route discovery if:

- o it receives both secure and unsecure DIOs or Secure DIOs with different Security Configurations pertaining to this route

discovery (i.e., referring to the same RPLInstanceID and DODAGID combination) prior to joining; or

- o it can not use the Security Configuration found in the Secure DIOs pertaining to this route discovery.

When a router (an Intermediate Router or a Target) joins a temporary DAG being created using Secure DIOs, it MUST remember the common Security Configuration used in the received Secured DIOs and MUST use this configuration to secure all the control messages (DIOs and P2P-DR0s) it generates.

If an Intermediate Router (or a Target) encounters a control message (a DIO or a P2P-DR0 or a P2P-DR0-ACK) pertaining to this route discovery that is either not secure or does not follow the Security Configuration the router remembers for this route discovery, the router MUST enter the "lock down" mode for the remainder of its stay in this temporary DAG. An Intermediate Router (or a Target) in the "lock down" mode MUST NOT generate or process any control message (irrespective of the Security Configuration used) pertaining to this route discovery. If the Origin receives a control message (a P2P-DR0) that does not follow the Security Configuration the Origin has chosen for this route discovery, it MUST discard the received message with no further processing.

[12.](#) Packet Forwarding Along a Route Discovered Using P2P-RPL

An Origin uses the Source Routing Header (SRH) [[RFC6554](#)] to send a packet along a Source Route discovered using P2P-RPL.

Travel along a Hop-by-hop Route, established using P2P-RPL, requires specifying the RPLInstanceID and the DODAGID (of the temporary DAG used for the route discovery) to identify the route. This is because a P2P-RPL route discovery does not use globally unique RPLInstanceID values and hence both the RPLInstanceID (a local value assigned by

the Origin) and the DODAGID (an IPv6 address of the Origin) are required to uniquely identify a P2P-RPL Hop-by-hop Route to a particular destination.

An Origin includes an RPL option [[RFC6553](#)] inside the IPv6 hop-by-hop options header of a packet to send it along a Hop-by-hop Route established using P2P-RPL. For this purpose, the Origin MUST set the DODAGID of the temporary DAG used for the route discovery as the source IPv6 address of the packet. Further, the Origin MUST specify inside the RPL option the RPLInstanceID of the temporary DAG used for the route discovery and set the 0 flag inside the RPL option to one. On receiving this packet, an Intermediate Router checks the 0 flag

and correctly infer the source IPv6 address of the packet as the DODAGID of the Hop-by-hop Route. The router then uses the DODAGID, the RPLInstanceID and the destination address to identify the routing state to be used to forward the packet further.

13. Interoperability with Core RPL

This section describes how RPL routers that implement P2P-RPL interact with RPL routers that do not. In general, P2P-RPL operation does not affect core RPL operation and vice versa. However, core RPL does allow a router to join a DAG as a leaf node even if it does not understand the Mode of Operation (MOP) used in the DAG. Thus, an RPL router that does not implement P2P-RPL may conceivably join a temporary DAG being created for a P2P-RPL route discovery as a leaf node and maintain its membership even though the DAG no longer exists. This may impose a drain on the router's memory. However, such RPL-only leaf nodes do not interfere with P2P-RPL route discovery since a leaf node may only generate a DIO advertising an INFINITE_RANK and all routers implementing P2P-RPL are required to discard such DIOs. Note that core RPL does not require a router to join a DAG whose MOP it does not understand. Moreover, RPL routers in a particular deployment may have strict restrictions on the DAGs they may join, thereby mitigating the problem.

The P2P-RPL mechanism described in this document works best when all the RPL routers in the LLN implement P2P-RPL. In general, the ability to discover routes as well as the quality of discovered routes would deteriorate with the fraction of RPL routers that

implement P2P-RPL.

14. Security Considerations

In general, the security considerations for the operation of P2P-RPL are similar to the ones for the operation of RPL (as described in [Section 19 of \[RFC6550\]](#)). Sections [6.1](#) and [10](#) of RPL specification [\[RFC6550\]](#) describe RPL's security framework that provides data confidentiality, authentication, replay protection and delay protection services. This security framework can also be used in P2P-RPL after taking in account the constraints specified in [Section 11](#). P2P-RPL requires all routers participating in a secure route discovery to use the Security Configuration decided by the Origin. The intention is to avoid compromising the overall security of a route discovery due to some routers using a weaker Security Configuration. With "lock down" mechanism, described in [Section 11](#), in effect, it is unlikely that an Origin would accept a route discovered under a Security Configuration other than the one it

intended. Any attempt to use a different Security Configuration (than the one the Origin intended) is likely to result, in the worst case, in the failure of the route discovery process. In the best case scenario, any such attempt by a rogue router would result in its neighbors entering the "lock down" mode and acting as firewalls to allow the route discovery to proceed in the remaining network.

RPL specification describes three modes of security: unsecured, pre-installed and authenticated. In the unsecured mode, secure control messages are not used and the only available security is the one provided by the link layer protocols. In the pre-installed mode, all the nodes use a pre-installed group key to join a secure DAG as the "routers" or "hosts", where the term "router" means a node that is capable of forwarding packets received from its parents or children in the DAG and the term "host" refers to nodes that can not function as "routers". In the authenticated mode, the nodes can join a secure DAG as "hosts" using the pre-installed key but then need to authenticate themselves to a key server to obtain the key that will allow them to work as "routers". The temporary DAG created for a P2P-RPL discovery can not be used for routing packets. Hence, it is not meaningful to say that a node joins this DAG as a "router" or a "host" in the sense defined above. Hence, in P2P-RPL, there is no

distinction between the pre-installed and authenticated modes. A router can join a temporary DAG created for a secure P2P-RPL route discovery only if it can support the Security Configuration in use, which also specifies the key in use. It does not matter whether the key is pre-installed or dynamically acquired. The router must have the key in use before it can join the DAG being created for a secure P2P-RPL route discovery.

If a rogue router can support the Security Configuration in use (in particular, it knows the key in use), it can join the secure P2P-RPL route discovery and cause a variety of damage. Such a rogue router could advertise false information in its DIOs in order to include itself in the discovered route(s). It could generate bogus P2P-DRO messages carrying bad routes or maliciously modify genuine P2P-DRO messages it receives. A rogue router acting as the Origin could launch denial of service attacks against the LLN deployment by initiating fake P2P-RPL route discoveries. Here, RPL's authenticated mode operation would be useful, where a node can obtain the key to use for a P2P-RPL route discovery only after proper authentication.

Since a P2P-DRO message travels along a Source Route specified inside the message, some of the security concerns that led to the deprecation of Type 0 routing header [RFC5095] may apply. To avoid the possibility of a P2P-DRO message traveling in a routing loop, this document requires each Intermediate Router to confirm that the Source Route listed inside the message does not contain any routing

loop involving itself before the router could forward the message further. As specified in [Section 9.6](#), this check involves the router making sure that its IPv6 addresses do not appear multiple times inside the Source Route with one or more other IPv6 addresses in between.

[15.](#) IANA Considerations

[15.1.](#) Additions to Mode of Operation

This document defines a new Mode of Operation, entitled "P2P Route Discovery Mode" (see [Section 6](#)), assigned a value TBD1 from the "Mode of Operation" space [to be removed upon publication: <http://www.iana.org/assignments/rpl/rpl.xml#mop>] [RFC6550]. IANA is

requested to allocate a suitable value to TBD1. The string TBD1 in this document should be replaced by the allocated value. The previous two sentences should be removed before publication.

Value	Description	Reference
TBD1	P2P Route Discovery Mode of Operation	This document

Mode of Operation

15.2. Additions to RPL Control Message Options

This document defines a new RPL option: "P2P Route Discovery Option" (see [Section 7](#)), assigned a value TBD2 from the "RPL Control Message Options" space [to be removed upon publication: <http://www.iana.org/assignments/rpl/rpl.xml#control-message-options>] [RFC6550]. IANA is requested to allocate a suitable value to TBD2. The string TBD2 in this document should be replaced by the allocated value. The previous two sentences should be removed before publication.

Value	Meaning	Reference
TBD2	P2P Route Discovery	This document

RPL Control Message Options

15.3. Additions to RPL Control Codes

This document defines the following new RPL messages:

- o "P2P Discovery Reply Object" (see [Section 8](#)), assigned a value TBD3 from the "RPL Control Codes" space [to be removed upon publication: <http://www.iana.org/assignments/rpl/rpl.xml#control-codes>]

[RFC6550]. IANA is requested to allocate TBD3 from the range 0x00-0x7F to indicate a message without security enabled. The string TBD3 in this document should be replaced by the allocated value. The previous two sentences should be removed before publication.

- o "Secure P2P Discovery Reply Object" (see [Section 8.1](#)), assigned a value TBD4 from the "RPL Control Codes" space [to be removed upon publication:
<http://www.iana.org/assignments/rpl/rpl.xml#control-codes>]
[RFC6550]. IANA is requested to allocate TBD4 from the range 0x80-0xFF to indicate a message with security enabled. The string TBD4 in this document should be replaced by the allocated value. The previous two sentences should be removed before publication.
- o "P2P Discovery Reply Object Acknowledgement" (see [Section 10](#)), assigned a value TBD5 from the "RPL Control Codes" space [to be removed upon publication:
<http://www.iana.org/assignments/rpl/rpl.xml#control-codes>]
[RFC6550]. IANA is requested to allocate TBD5 from the range 0x00-0x7F to indicate a message without security enabled. The string TBD5 in this document should be replaced by the allocated value. The previous two sentences should be removed before publication.
- o "Secure P2P Discovery Reply Object Acknowledgement" (see [Section 10](#)), assigned a value TBD6 from the "RPL Control Codes" space [to be removed upon publication:
<http://www.iana.org/assignments/rpl/rpl.xml#control-codes>]
[RFC6550]. IANA is requested to allocate TBD6 from the range 0x80-0xFF to indicate a message with security enabled. The string TBD6 in this document should be replaced by the allocated value. The previous two sentences should be removed before publication.

Code	Description	Reference
TBD3	P2P Discovery Reply Object	This document
TBD4	Secure P2P Discovery Reply Object	This document
TBD5	P2P Discovery Reply Object Acknowledgement	This document
TBD6	Secure P2P Discovery Reply Object Acknowledgement	This document

RPL Control Codes

[16.](#) Acknowledgements

Authors gratefully acknowledge the contributions of the following individuals (in alphabetical order) in the development of this document: Dominique Barthel, Jakob Buron, Cedric Chauvenet, Thomas Clausen, Robert Cragie, Ralph Droms, Adrian Farrel, Stephen Farrell, Brian Haberman, Ted Humpal, Richard Kelsey, Phil Levis, Charles Perkins, Joseph Reddy, Michael Richardson, Zach Shelby, Martin Stiernerling, Pascal Thubert, Hristo Valev and JP Vasseur.

[17.](#) References

[17.1.](#) Normative References

[I-D.ietf-roll-terminology]

Vasseur, J., "Terminology in Low power And Lossy Networks", [draft-ietf-roll-terminology-10](#) (work in progress), January 2013.

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

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.

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

[RFC6550] Winter, T., Thubert, P., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R.

Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", [RFC 6550](#), March 2012.

- [RFC6551] Vasseur, JP., Kim, M., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", [RFC 6551](#), March 2012.

[17.2.](#) Informative References

- [I-D.ietf-roll-p2p-measurement]
Goyal, M., Baccelli, E., Brandt, A., and J. Martocci, "A Mechanism to Measure the Routing Metrics along a Point-to-point Route in a Low Power and Lossy Network", [draft-ietf-roll-p2p-measurement-09](#) (work in progress), February 2013.
- [RFC5095] Abley, J., Savola, P., and G. Neville-Neil, "Deprecation of Type 0 Routing Headers in IPv6", [RFC 5095](#), December 2007.
- [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.
- [RFC6552] Thubert, P., "Objective Function Zero for the Routing Protocol for Low-Power and Lossy Networks (RPL)", [RFC 6552](#), March 2012.
- [RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", [RFC 6553](#), March 2012.
- [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](#), March 2012.

Internet-Draft

[draft-ietf-roll-p2p-rpl-17](#)

March 2013

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

Phone: +33-169-335-511
Email: Matthias.Philipp@inria.fr

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

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

Jerald Martocci
Johnson Controls

507 E Michigan St
Milwaukee, WI 53202
USA

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

Goyal, et al.

Expires September 21, 2013

[Page 39]