

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: April 29, 2011

M. Goyal, Ed.
University of Wisconsin Milwaukee
E. Baccelli, Ed.
INRIA
P2P. Team
October 26, 2010

A Mechanism to Measure the Quality of a Point-to-point Route in a Low
Power and Lossy Network
draft-goyal-roll-p2p-measurement-01

Abstract

This document specifies a mechanism that enables an RPL node to measure the quality of an existing route to/from another RPL node in a low power and lossy network, thereby allowing the node to decide if it wants to initiate the discovery of a more optimal route.

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 April 29, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<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	3
1.1. Terminology	4
2. Functional Overview	4
3. The Measurement Object (MO)	5
4. Originating an MO To Measure a P2P Route	6
4.1. From the Origin Node to the Target Node	7
4.2. From the Target Node to the Origin Node	8
5. Processing a Received MO at an Intermediate Router	8
6. Processing a Received MO at the Target Node	9
7. Processing a Received MO at the Origin Node	11
8. Security Considerations	11
9. IANA Considerations	11
10. Authors and Contributors	11
11. References	12
11.1. Normative References	12
11.2. Informative References	12
Authors' Addresses	13

1. Introduction

Point to point (P2P) communication between arbitrary nodes in a Low power and Lossy Network (LLN) is a key requirement for many applications [RFC5826][RFC5867]. RPL [I-D.ietf-roll-rpl], the IPv6 Routing Protocol for LLNs, constrains the LLN topology to a Directed Acyclic Graph (DAG) built to optimize routing costs to reach the DAG's root and requires the P2P routes to use the DAG links only. Such P2P routes may potentially be suboptimal and may lead to traffic congestion near the DAG root. Additionally, RPL is a proactive routing protocol and hence all P2P routes must be established ahead of the time they are used.

To ameliorate situations, where RPL's P2P routing functionality does not meet the requirements, [I-D.ietf-roll-p2p-rpl] describes a reactive mechanism to discover P2P routes that meet the specified performance characteristics. This mechanism, henceforth referred to as the reactive P2P route discovery, requires the specification of "good enough criteria", in terms of constraints on aggregated values of the relevant routing metrics [I-D.ietf-roll-routing-metrics], that the discovered routes must satisfy. In some cases, the application requirements or the LLN's topological features allow a node to infer the good enough criteria intrinsically. For example, the application may require the end-to-end loss rate and/or latency on the route to be below certain thresholds or the LLN topology may be such that a router can safely assume its destination to be less than a certain number of hops away from itself.

When the existing P2P routes are deemed unsatisfactory by the application layer but the node does not intrinsically know the good enough criteria, it may be necessary for the node to determine the aggregated values of relevant routing metrics along the existing routes. This knowledge will allow the node to frame a reasonable good enough criteria and initiate a reactive P2P route discovery to determine better routes. For example, if the router determines the aggregate ETX [I-D.ietf-roll-routing-metrics] along an existing route to be "x", it can use " $ETX < x*y$ ", where y is a certain fraction, as a constraint in the good enough criteria. Note that it is important that the good enough criteria is not overly strict; otherwise the route discovery may fail even though routes, much better than the ones being currently used, exist.

This document specifies a mechanism that enables an RPL node to measure the aggregated values of the routing metrics along an existing route to/from another RPL node in an LLN, thereby allowing the node to decide if it wants to initiate the reactive discovery of a more optimal route and determine the good enough criteria to be used for this purpose.

1.1. 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], [I-D.ietf-roll-rpl] and [I-D.ietf-roll-p2p-rpl]. Specifically, the term node refers to an RPL router or an RPL host as defined in [I-D.ietf-roll-rpl]. The following terms, originally defined in [I-D.ietf-roll-p2p-rpl], are redefined in the following manner.

Origin Node: The origin node refers to the node that initiates the measurement process defined in this document and is one end point of the P2P route being measured.

Target Node: The target node refers to the other end of the P2P route being measured.

Intermediate Router: A router, other than the origin and the target node, on the P2P route being measured.

2. Functional Overview

The mechanism described in this document can be used by an origin node to measure the aggregated values of the routing metrics along a P2P route to/from a target node in the LLN. Such a route could be a source route or a hop-by-hop route established using RPL [I-D.ietf-roll-rpl] or the reactive P2P route discovery [I-D.ietf-roll-p2p-rpl].

When an origin node desires to measure the aggregated values of the routing metrics along a P2P route from itself to a target node, it sends a Measurement Request message along that route. The Measurement Request message accumulates the values of the relevant routing metrics as it travels towards the target node. Upon receiving the Measurement Request message, the target node unicasts a Measurement Reply message, carrying the accumulated values of the routing metrics, back to the origin node.

When an origin node desires to measure the aggregated values of the routing metrics along a P2P route from a target node to itself, it unicasts a Measurement Request message, specifying the routing metrics to be measured, to the target node. On receiving the Measurement Request message, the target node sends a Measurement

Reply message to the origin node along the P2P route to be measured. The Measurement Reply message accumulates the values of the relevant routing metrics as it travels towards the origin node.

3. The Measurement Object (MO)

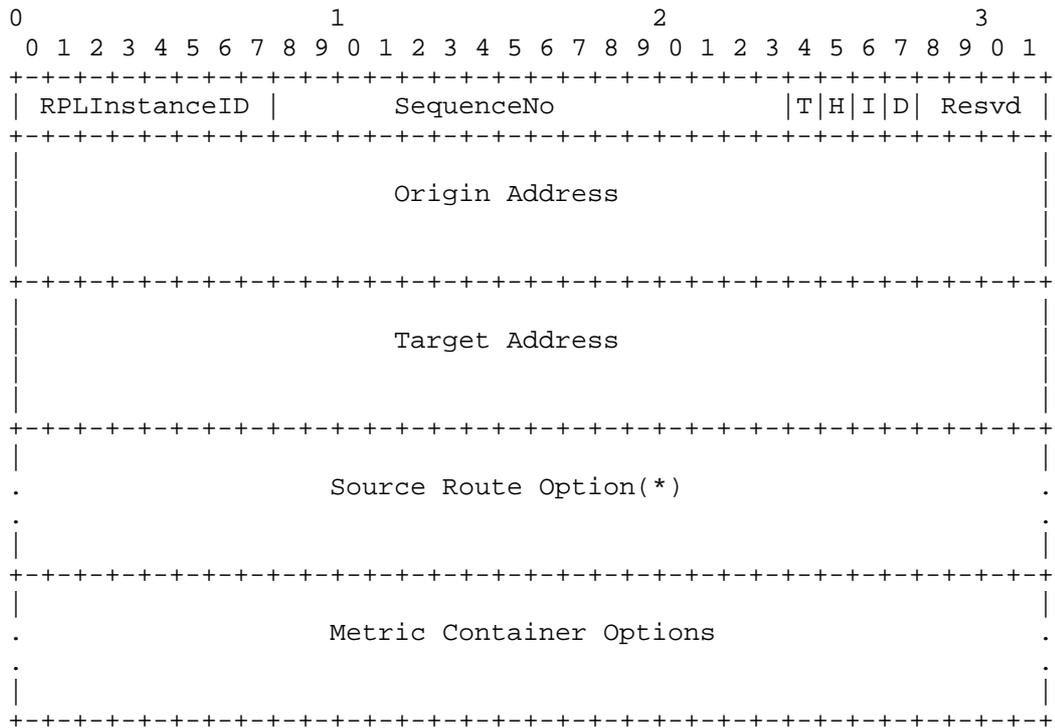


Figure 1: Format of the Measurement Object (MO)

This document defines a new RPL Control Message type, the Measurement Object (MO) with code 0x05 (to be confirmed by IANA) that serves as both Measurement Request and Measurement Reply. The format of an MO is shown in Figure 1. An MO consists of the following fields:

- o RPLInstanceID: Relevant only if the MO travels along a hop-by-hop route. This field identifies the RPLInstanceID of the hop-by-hop route.
- o SequenceNo: A 16-bit sequence number that uniquely identifies a Measurement Request and the corresponding Measurement Reply to the origin node.

- o T: The type flag. This flag is set if the MO represents a Measurement Request. The flag is cleared if the MO is a Measurement Reply.
- o H: This flag is set if the MO travels along a hop-by-hop route. In that case, the hop-by-hop route is identified by the RPLInstanceID and, if required, the Origin/Target Address serving as the DODAGID. This flag is cleared if the MO travels along a source route. In that case, the MO MUST contain a Source Route option [I-D.ietf-roll-p2p-rpl]. Note that, in case of a P2P route along a non-storing DAG, it is possible that an MO message travels along a hop-by-hop route till the DAG's root, which then sends it along a source route to its destination. In that case, the DAG root will reset the H flag and also insert a Source Route option in the MO.
- o I: A flag that indicates which of the two - the Origin Address and the Target Address - indicates the DODAGID for the hop-by-hop route. This flag is relevant only if the MO travels along a hop-by-hop route (i.e., H flag is set) and a local RPLInstanceID has been specified to identify the hop-by-hop route. This flag is set if the Origin Address indicates the DODAGID; the flag is cleared if the Target Address indicates the DODAGID.
- o D: A flag that indicates the direction of the P2P route. This flag is set when the route to be measured is from the origin node to the target node. Otherwise, the flag is cleared.
- 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 Origin Address: The IPv6 address of the origin node.
- o Target Address: The IPv6 address of the target node.
- o Source Route Option: An MO MUST contain one Source Route option if it travels along a source route.
- o Metric Container Options: An MO MUST contain one or more Metric Container options to carry the routing metric objects [I-D.ietf-roll-routing-metrics].

4. Originating an MO To Measure a P2P Route

4.1. From the Origin Node to the Target Node

If the origin node intends to measure the routing metric values along a P2P route towards a target node, it generates an MO message and sets its fields as follows:

- o RPLInstanceID: If the P2P route is a hop-by-hop route, the origin node specifies the RPLInstanceID to identify the route in this field. This field is not relevant if the P2P route is a source route specified in the Source Route option. This document RECOMMENDS a value 10000000 for this field if the P2P route is a source route.
- o SequenceNo: The origin node assigns a sequence number to the MO to uniquely identify the corresponding Measurement Reply.
- o T: The T flag is set to indicate that MO represents a Measurement Request.
- o H: The H flag is set if the MO travels along a hop-by-hop route.
- o I: This field is relevant only if the H flag is set and the RPLInstanceID is a local value. The origin node sets this flag if the Origin Address indicates the DODAGID. The origin node clears this flag if the Target Address indicates the DODAGID.
- o D: This flag is set.
- o Origin Address, Target Address: These fields are set to the IPv6 addresses of the origin and target nodes respectively. If the H flag is set and the RPLInstanceID is a local value, the Origin Address or the Target Address MUST also indicate the DODAGID value required to identify the hop-by-hop route.
- o Source Route Option: If the P2P route is a source route (i.e., the H flag is cleared), the Source Route option MUST be present and MUST include a complete source route to the target node in forward direction (excluding the addresses of the origin and target nodes).
- o Metric Container Options: The origin node MUST also include one or more Metric Container options containing relevant routing metric objects to accumulate the costs for these metrics along the P2P route. The origin node also initiates the routing metric objects by including the local values of the routing metrics for the first hop on the P2P route.

After setting the MO fields as described above, the origin node MUST

unicast the MO message to the next hop on the P2P route. The origin node MAY include a Record Route IPv6 Extension Header, proposed in [I-D.thubert-6man-reverse-routing-header], in the MO message to accumulate a reverse route that the target node can use to send the Measurement Reply back to the origin node.

4.2. From the Target Node to the Origin Node

If the origin node intends to measure the routing metric values along a P2P route from a target node to itself, it generates an MO message and sets its fields as follows:

- o SequenceNo: The origin node assigns a sequence number to the MO to uniquely identify the corresponding Measurement Reply.
- o T: The T flag is set to indicate that MO represents a Measurement Request.
- o D: This flag is cleared.
- o Origin Address, Target Address: These fields are set to the IPv6 addresses of the origin and target nodes respectively.
- o Source Route Option: In this case, the MO SHOULD NOT include any Source Route option.
- o Metric Container Options: The origin node MUST include one or more Metric Container options containing relevant routing metric objects to accumulate the costs for these metrics along the P2P route. These routing metric objects MUST be empty.

The other fields in the MO are not relevant in this case and SHOULD be set to zero. After setting the MO fields as described above, the origin node MUST unicast the MO message to the target node.

5. Processing a Received MO at an Intermediate Router

When a node receives an MO, it examines if one of its IPv6 addresses is listed as the Origin Address or the Target Address. If not, the node checks if H bit is clear (i.e., the MO is traveling along a source route). If yes, the node checks the Address[0] field inside the Source Route Option contained in the MO. The node MUST drop the MO with no further processing and send an ICMPv6 Destination Unreachable error message to the source of the message (the Origin Address if the MO is a Measurement Request; otherwise the Target Address) if the received MO has a clear H bit but does not contain a Source Route Option or if the Address[0] inside the Source Route

option does not match one of the node's IPv6 address.

The node then determines the next hop on the P2P route being measured. In case the received MO has a clear H flag, the Address[1] field (the second element in the Address vector) inside the Source Route Option is taken as the next hop. If the Source Route Option does not contain Address[1] element, the node checks the T flag inside the MO. If T flag is set, i.e., MO is a Measurement Request, the Target Address is taken as the next hop; otherwise the Origin Address is the next hop. If the received MO has H flag set, the node uses the RPLInstanceID, the ultimate destination of the MO (Target Address if T flag is set; otherwise the Origin Address) and, if RPLInstanceID is a local value, the DODAGID (the Origin Address if I flag is set; otherwise the Target Address) to determine the next hop for the MO. If the H flag in the MO is set and the node is the root of a non-storing DAG, indicated by the RPLInstanceID, the node MAY reset the H flag and insert a Source Route option in the MO to indicate a source route along which the MO should travel on rest of its way to its destination. The node MUST drop the MO with no further processing and send an ICMPv6 Destination Unreachable error message to the source of the message if it can not determine the next hop for the message.

After determining the next hop, the node updates the routing metric objects, contained in the Metric Container options inside the MO, either by updating the aggregated value for the routing metric or by attaching the local values for the metric inside the object. The node MUST drop the MO with no further processing and send a suitable ICMPv6 error message to the source of the message if the node does not know the relevant routing metric values for the next hop.

After updating the routing metrics, the node MUST unicast the MO to the next hop. If the MO to be forwarded has a clear H flag, the node MUST ensure that the Address vector in the Source Route option contains the next hop address as the first element.

6. Processing a Received MO at the Target Node

When a node receives an MO, it examines if one of its IPv6 addresses is listed as the Target Address. If yes, the node checks the T flag. The node MUST drop the MO with no further processing and optionally log an error if the T flag is clear (i.e. the received MO is a Measurement Reply).

The target node then checks the D flag to determine the direction of the P2P route to be measured. If the D flag is set (i.e., the P2P route to be measured is from the origin node to the target node), the

target node updates the routing metrics objects in the Metric Container options if required, removes the Source Route Option if present and clears the T bit thereby converting the MO into a Measurement Reply. The target node then unicasts the updated MO back to the origin node. For this purpose, the target node MAY use the reverse route accumulated in the Record Route IPv6 Extension Header [I-D.thubert-6man-reverse-routing-header] if present in the received MO message.

If the D flag in the received MO message is clear (i.e., the P2P route to be measured is from the target node to the origin node), the target node selects the P2P route to be measured and modifies the following MO fields:

- o RPLInstanceID: If the P2P route is a hop-by-hop route, the target node specifies in this field the RPLInstanceID associated with the route. This field is not relevant if the P2P route is a source route. This document RECOMMENDS a value 10000000 for this field if the P2P route is a source route.
- o T: The T flag is cleared to indicate that MO represents a Measurement Reply.
- o H: The H flag is set if the P2P route is a hop-by-hop one.
- o I: If the H flag is set and the RPLInstanceID is a local value, the target node sets this flag if the Origin Address indicates the DODAGID. The target node clears this flag if the Target Address indicates the DODAGID.
- o D: This flag is cleared.
- o Source Route Option: If the P2P route is a source route, the Source Route option MUST be present and MUST include a complete source route from the target node to the origin node (excluding the addresses of the target and origin nodes).
- o Metric Container Options: The target node MUST initiate the routing metric objects inside the Metric Container options by including the local values of the routing metrics for the first hop on the P2P route.

The target node need not modify the other fields in the received MO. After these modifications, the target node MUST unicast the MO message to the next hop on the P2P route.

7. Processing a Received MO at the Origin Node

When a node receives an MO, it examines if one of its IPv6 addresses is listed as the Origin Address. If yes, the node checks the T flag. The node MUST drop the MO with no further processing and optionally log an error if the T flag is set (i.e. the received MO is a Measurement Request) or if the node has no recollection of sending a Measurement Request with the sequence number listed in the received MO.

If the D flag in the received MO is clear (i.e., the P2P route to be measured is from the target node to the origin node), the origin node MUST update the routing metrics objects in the Metric Container options if required.

The origin node can now examine the routing metric objects inside the Metric Container options to evaluate the quality of the measured P2P route. If a routing metric object contains local metric values recorded by enroute nodes, the origin node MAY aggregate these local values into an end-to-end value as per the aggregation rules for the metric.

8. Security Considerations

TBA

9. IANA Considerations

TBA

10. Authors and Contributors

In addition to the editors, the authors of this document include the following individuals (listed in alphabetical order).

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, Wakefieldm WF4 4WA, UK. Phone: +44 1924910888; Email: robert.cragie@gridmerge.com

Jerald Martocci, Johnson Controls, Milwaukee, WI 53202, USA. Phone: +1 414 524 4010; Email: jerald.p.martocci@jci.com

Charles Perkins, Tellabs Inc., USA. Email: charliep@computer.org

Authors gratefully acknowledge the contributions of Richard Kelsey and Zach Shelby in the development of this document.

11. References

11.1. Normative References

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

11.2. Informative References

[I-D.ietf-roll-p2p-rpl]
Goyal, M. and E. Baccelli, "Reactive Discovery of Point-to-Point Routes in Low Power and Lossy Networks", draft-ietf-roll-p2p-rpl-00 (work in progress), August 2010.

[I-D.ietf-roll-routing-metrics]
Vasseur, J., Kim, M., Networks, D., Dejean, N., and D. Barthel, "Routing Metrics used for Path Calculation in Low Power and Lossy Networks", draft-ietf-roll-routing-metrics-11 (work in progress), October 2010.

[I-D.ietf-roll-rpl]
Winter, T., Thubert, P., Brandt, A., Clausen, T., Hui, J., Kelsey, R., Levis, P., Networks, D., Struik, R., and J. Vasseur, "RPL: IPv6 Routing Protocol for Low power and Lossy Networks", draft-ietf-roll-rpl-13 (work in progress), October 2010.

[I-D.ietf-roll-terminology]
Vasseur, J., "Terminology in Low power And Lossy Networks", draft-ietf-roll-terminology-04 (work in progress), September 2010.

[I-D.thubert-6man-reverse-routing-header]
Thubert, P., "Reverse Routing Header", draft-thubert-6man-reverse-routing-header-00 (work in progress), June 2010.

[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 53211
USA

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

Emmanuel Baccelli (editor)
INRIA

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

P2P Team

