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

M. Goyal, Ed.
University of Wisconsin
Milwaukee
E. Baccelli, Ed.
INRIA
A. Brandt
Sigma Designs
R. Cragie
Gridmerge Ltd
J. Martocci
Johnson Controls
July 11, 2011

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

Abstract

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

Table of Contents

1. Introduction

Point to point (P2P) communication between arbitrary routers 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 criteria. This mechanism, henceforth referred to as the reactive P2P route discovery, requires the specification of routing constraints [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 router to infer the routing constraints 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 routes are deemed unsatisfactory but the router does not intrinsically know the routing constraints to be used in P2P route discovery, it may be necessary for the router to determine the aggregated values of the routing metrics along the existing route. This knowledge will allow the router to frame reasonable routing constraints for use in P2P route discovery to determine a better route. 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 the routing constraint for use in P2P route discovery. Note that it is important that the routing constraints are not overly strict; otherwise the P2P route discovery may fail even though a route, much better than the one currently being used, exists.

This document specifies a mechanism that enables an RPL router to measure the aggregated values of the routing metrics along an existing route to another RPL router in an LLN, thereby allowing the router to decide if it wants to initiate the reactive discovery of a more optimal route and determine the routing constraints 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]. The following terms, originally defined in [I-D.ietf-roll-p2p-rpl], are redefined in the following manner.

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

Target: The target refers to the router at the end point of the P2P route being measured.

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

2. Functional Overview

The mechanism described in this document can be used by an origin to measure the aggregated values of the routing metrics along a P2P route to a target 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]. The origin sends a Measurement Request message along the route. The Measurement Request accumulates the values of the routing metrics as it travels towards the target. Upon receiving the Measurement Request, the target unicasts a Measurement Reply message, carrying the accumulated values of the routing metrics, back to the origin.

3. The Measurement Object (MO)

```
 0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 | RPLInstanceID |  SequenceNo   | Compr |T|H|A|R|  Num  | Index |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 |                      Origin Address                           |
 |                                                               |
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 |                      Target Address                           |
 |                                                               |
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 .                      Address[1..Num]                          .
 .                                                               .
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 .                  Metric Container Option(s)                   .
 .                                                               .
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 1: Format of the Measurement Object (MO)

This document defines a new RPL Control Message type, the Measurement Object (MO), with code 0x06 (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:

- RPLInstanceID: Relevant only if the MO travels along a hop-by-hop route. This field identifies the RPLInstanceID of the hop-by-hop route being measured. If the route being measured is a source route, this field MUST be set to 10000000 on transmission and ignored on reception.

- SequenceNo: An 8-bit sequence number that uniquely identifies a Measurement Request and the corresponding Measurement Reply.

- Compr: A 4-bit unsigned integer indicating the number of prefix octets that are elided from the IPv6 addresses in Origin/Target Address fields and the Address vector. For example, Compr value

will be 0 if full IPv6 addresses are carried in the Origin/Target Address fields and the Address vector.

- Type (T): This flag is set if the MO represents a Measurement Request. The flag is cleared if the MO is a Measurement Reply.

- Hop-by-hop (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 the RPLInstanceID is a local value, the Origin Address serving as the DODAGID. This flag is cleared if the MO travels along a source route specified in the Address vector. Note that, in case the P2P route being measured lies along a non-storing DAG, an MO message may travel along a hop-by-hop route till it reaches 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 the source route to the destination inside the Address vector.

- Accumulate Route (A): This flag is relevant only if the MO represents a Measurement Request that travels along a hop-by-hop route represented by a local RPLInstanceID. When this flag is relevant, a value 1 in the flag indicates that the Measurement Request MUST accumulate a source route for use by the target to send the Measurement Reply back to the origin. In this case, the intermediate routers MUST add their IPv6 addresses (after eliding Compr number of prefix octets) to the Address vector in the manner specified later.

- Reverse (R): This flag is relevant only if the MO represents a Measurement Request that travels along a source route, specified in the Address vector, to the target. When this flag is relevant, a value 1 in the flag indicates that the Address vector contains a complete source route from the origin to the target, which can be used, after reversal, by the target to source route the Measurement Reply message back to the origin.

- Num: This field indicates the number of fields in the Address vector. If the value of this field is zero, the Address vector is not present in the MO.

- Index: If the Measurement Request is traveling along a source route contained in the Address vector, this field indicates the index in the Address vector of the next hop on the route. If the Measurement Request is traveling along a hop-by-hop route with a local RPLInstanceID and the A flag is set, this field indicates the index in the Address vector where an intermediate router receiving the MO message must store its IPv6 address. Otherwise, this field MUST be set to zero on transmission and ignored on

reception.

o Origin Address: An IPv6 address of the origin after eliding Compr number of prefix octets. If the MO is traveling along a hop-by-hop route and the RPLInstanceID field indicates a local value, the Origin Address field MUST contain the DODAGID value that, along with the RPLInstanceID, uniquely identifies the hop-by-hop route being measured.

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

o Address[1..Num]: A vector of IPv6 addresses (with Compr number of prefix octets elided) representing a (partial) route from the origin to the target:

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

   * The total number of elements inside the Address vector is given by the Num field.

   * When the Measurement Request is traveling along a hop-by-hop route with local RPLInstanceID and has the A flag set, the Address vector is used to accumulate a route to be used by the target to send the Measurement Reply back to the origin. In this case, the route MUST be accumulated in the forward direction, i.e., from the origin to the target. The target router would reverse this route to obtain a source route from itself to the origin. The IPv6 addresses in the accumulated route MUST be accessible in the backward direction. An intermediate router adding its address to the Address vector MUST ensure that its address does not already exist in the vector.

   * When the Measurement Request is traveling along a source route, the Address vector MUST contain a complete route to the target and the IPv6 addresses in the Address vector MUST be accessible in the forward direction, i.e., from the origin to the target. A router (the origin or an intermediate router) specifying a route to the target in the Address vector MUST ensure that the vector does not contain any address more than once. The origin may set the R flag in the MO if the route in the Address vector represents a complete route from the origin to the target and this route can be used after reversal by the target to send the Measurement Reply message back to the origin.

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

      * The Address vector MUST NOT contain any multicast addresses.

- Metric Container Options: An MO MUST contain one or more Metric Container options to accumulate routing metric values for the route being measured.

4. Originating a Measurement Request

If an origin needs to measure the routing metric values along a P2P route towards a target, it generates an MO message and sets its fields in the manner described above. Specifically, the origin MUST set the T flag to 1 to indicate that the MO represents a Measurement Request.

If a source route is being measured, the origin MUST do the following:

- specify the complete source route to the target inside the Address vector;

- specify in the Num field the number of address elements in the Address vector;

- set the Index field to value zero;

- set the R flag if the route in the Address vector can be used after reversal by the target to source route the Measurement Reply message back to the origin.

If a hop-by-hop route with a local RPLInstanceID is being measured and the origin desires the MO to accumulate a source route for the target to send the Measurement Reply message back, it MUST do the following:

- set A flag to 1;

- include a suitably sized, empty Address vector (with all bits set to zero) in the MO;

- specify in the Num field the number of address elements that can fit inside the Address vector;

- set the Index field to value zero.

The origin MUST include one or more Metric Container options inside the MO that carry the routing metric objects of interest. If required, the origin must also initiate these routing metric objects

by including the values of the routing metrics for the first hop on the P2P route being measured.

After setting the MO fields as described above, the origin MUST unicast the MO message to the next hop on the P2P route.

5. Processing a Measurement Request at an Intermediate Router

When a router receives an MO, it examines if one of its IPv6 addresses is listed as the Origin or the Target Address. If not, the router processes the received message in the following manner.

An intermediate router MUST discard the packet with no further processing if the received MO is not a Measurement Request.

The router then determines the next hop on the P2P route being measured. In case the received MO has a clear H flag, the router increments the Index field and uses the Address[Index] element as the next hop. If this element does not exist, the router uses the Target Address as the next hop.

If the received MO has H flag set to 1, the router uses the RPLInstanceID, the Target Address and, if RPLInstanceID is a local value, the DODAGID (same as the Origin Address) to determine the next hop for the MO. Also,

- o If the RPLInstanceID of the hop-by-hop route is a local value and the A flag is set, the router MUST store one of its IPv6 addresses (after eliding Compr bytes and making sure that the Address vector does not already contains one of its IPv6 addresses) at location Address[Index] and then increments the Index field.

- o If the router is the root of the non-storing DAG along which the received MO message has been traveling, the router MUST do the following:

  * reset the H, A and R flags;

  * insert a source route to the target inside the Address vector;

  * specify in the Num field the number of address elements in the Address vector;

  * set the Index field to value zero;

The router 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 router 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 router MUST drop the MO with no further processing and send a suitable ICMPv6 error message to the source of the message if the router does not know the relevant routing metric values for the next hop.

After updating the routing metrics, the router MUST unicast the MO to the next hop.

6. Processing a Measurement Request at the Target

When a router receives an MO, it examines if one of its IPv6 addresses is listed as the Target Address. If yes, the router processes the received message in the following manner.

An intermediate router MUST discard the packet with no further processing if the received MO is not a Measurement Request.

The target then updates the routing metrics objects in the Metric Container options if required and generates a Measurement Reply message. The received Measurement Request message can be trivially converted into the Measurement Reply by reseting the T flag to zero. The target MAY remove the Address vector from the Measurement Reply if desired. The target then unicasts the Measurement Reply back to the origin:

- o If the Measurement Request traveled along a DAG with a global RPLInstanceID, the Measurement Reply MAY be unicast back to the origin along the same DAG.

- o If the Measurement Request traveled along a hop-by-hop route with a local RPLInstanceID and the A flag inside the received message is set, the target MAY reverse the source route contained in the Address vector and use it to send the Measurement Reply back to the origin.

- o If the Measurement Request traveled along a source route and the R flag inside the received message it set, the target MAY reverse the source route contained in the Address vector and use it to send the Measurement Reply back to the origin.

7. Processing a Measurement Reply at the Origin

When a router receives an MO, it examines if one of its IPv6 addresses is listed as the Origin Address. If yes, the router processes the received message in the following manner.

The origin MUST discard the packet with no further processing if the received MO is not a Measurement Reply or if the origin has no recollection of sending a Measurement Request with the sequence number listed in the received MO.

The origin then examines 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 routers, the origin 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., Baccelli, E., Brandt, A., Cragie, R., and J. Martocci, "Reactive Discovery of Point-to-Point Routes in Low Power and Lossy Networks", draft-ietf-roll-p2p-rpl-03 (work in progress), May 2011.

[I-D.ietf-roll-routing-metrics]
Vasseur, J., Kim, M., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics used for Path Calculation in Low Power and Lossy Networks", draft-ietf-roll-routing-metrics-19 (work in progress), March 2011.

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

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

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

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

Authors' Addresses

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

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

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

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

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

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

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