

Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: March 8, 2012

M. Goyal
University of Wisconsin
Milwaukee
N. Dejean
Elster SAS
D. Barthel
France Telecom Orange
E. Baccelli
INRIA
J. Martocci
Johnson Controls
September 5, 2011

DIS Modifications
draft-goyal-roll-dis-modifications-00

Abstract

This document specifies the DIS flags and options that allow an RPL node to control how neighbor RPL routers respond to its solicitation for DIOs.

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 March 8, 2012.

Copyright Notice

Copyright (c) 2011 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
2. Terminology	4
3. New Flags in the DIS Base Object	4
4. DIS Options	5
4.1. Metric Container	5
4.2. Response Spreading Option	5
5. Applications	6
5.1. A Leaf Node Joining a DAG	6
5.2. Identifying A Defunct DAG	6
6. IANA Considerations	8
6.1. DIS Flags	8
6.2. RPL Control Message Options	9
7. Security Considerations	9
8. References	9
8.1. Normative References	9
8.2. Informative References	9
Authors' Addresses	10

1. Introduction

An RPL node can use a DODAG Information Solicitation (DIS) message to solicit DODAG Information Object (DIO) messages from its neighbor RPL routers. A DIS may carry a Solicited Information option that specifies the predicates of the DAG(s) the node is interested in. In the absence of a Solicited Information option, it is assumed that the node generating the DIS is interested in receiving DIOs for all the DAGs. A DIS can be multicast to all the in-range routers or it can be unicast to a specific neighbor router. RPL requires a router to consider the receipt of a multicast DIS as an inconsistency and hence reset its Trickle timers [RFC6206] for the matching DAGs. The receipt of a unicast DIS causes an RPL router to generate the DIOs for all the matching DAGs without resetting the Trickle timers.

Consider an RPL leaf node that desires to join a certain DAG. This node can either wait for its neighbor RPL routers to voluntarily transmit DIOs or it can proactively solicit DIOs using a DIS message. Voluntary DIO transmissions may happen after a very long time if the network is stable and the Trickle timer intervals have reached large values. Thus, proactively seeking DIOs using a DIS may be the only reasonable option. Since the node does not know which neighbor routers belong to the DAG, it must solicit the DIOs using a multicast DIS (with predicates of the desired DAG specified inside a Solicited Information option). On receiving this DIS, the neighbor routers that belong to the desired DAG will reset their Trickle timers and quickly transmit their DIOs. The downside of resetting Trickle timers is that the routers would continue to transmit the DIOs frequently for a considerable time interval. These DIO transmissions are unnecessary, consume precious energy and may contribute to congestion in the network.

There are other scenarios where resetting of Trickle timer following the receipt of a multicast DIS is not appropriate. For example, consider an RPL router that desires to free up memory by deleting state for the defunct DAGs it belongs to. Identifying a defunct DAG may require the node to solicit DIOs from its DAG parents using a multicast DIS.

To deal with the situations described above, this document specifies the DIS flags and options that allow an RPL node to control how neighbor RPL routers respond to its solicitation for DIOs:

- o Using routing constraints to limit the number of responding routers;
- o Whether the responding routers should reset their Trickle timers;

- o Whether the responding routers should send a unicast DIO or a multicast one;
- o The time interval over which the responding routers must schedule their DIO transmissions.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Additionally, this document uses terminology from [I-D.ietf-roll-terminology] and [I-D.ietf-roll-rpl]. Specifically, the term RPL node refers to an RPL router or an RPL host as defined in [I-D.ietf-roll-rpl].

3. New Flags in the DIS Base Object




Figure 1: Modified DIS Base Object

This document defines two new flags inside the DIS base object:

- o "No Inconsistency" (N) flag: On receiving a unicast/multicast DIS with N flag set, an RPL router MUST NOT reset the Trickle timers for the matching DAGs. Also, a DIO generated in response to a DIS with N flag set MUST always contain a Configuration option.
- o "DIO Type" (T) flag: This flag specifies whether the responding routers should transmit a multicast DIO or a unicast one. The responding router MUST transmit a multicast DIO if this flag is set.

The modified DIS base object is shown in Figure 1.

4. DIS Options

4.1. Metric Container

In order to limit the number of routers that will respond to a multicast DIS, this document allows the specification of routing constraints inside a DIS that a router must satisfy in order to respond to the DIS. These routing constraints are specified inside a Metric Container option contained in the DIS. Thus, this document allows the inclusion of a Metric Container option inside a DIS. An RPL router that receives a DIS with a Metric Container option MUST ignore any Metric object in it, and MUST evaluate the "mandatory" Constraint objects in it by comparing the constraint value to the aggregated value of the corresponding routing metric that the router maintains for the matching DAG(s). The aggregated routing metric values MUST satisfy all the mandatory constraints in order for the router to generate DIOs for the matching DAG(s).

4.2. Response Spreading Option




Figure 2: The Response Spreading Option

A multicast DIS may lead to a large number of RPL routers responding with DIOs. Concurrent transmissions by multiple routers are not desirable. Such transmissions may end up in collisions. Unicast DIOs may be able to avail of link-level retransmissions. However, multicast DIOs have no such protection. These transmissions and retransmissions may also cause congestion in the network. To avoid such problems, this document specifies an optional DIO response spreading mechanism.

This document defines a new RPL control message option called "Response Spreading", shown in Figure 2, with a recommended Type value 0x0A (to be confirmed by IANA). A Response Spreading option may be included only inside a multicast DIS message. An RPL router that responds to a multicast DIS, that includes a Response Spreading option, MUST wait for a time uniformly chosen in the interval $[0..2^{\text{SpreadingInterval}}]$, expressed in ms, before attempting to transmit its DIO. If the DIS does not include a Response Spreading option, the node is free to transmit the DIO as it otherwise would.

5. Applications

This section details two example mechanisms that use the DIS flags and options defined in this document. The first mechanism describes how a leaf node may join a desired DAG. The second mechanism details how a node may identify defunct DAGs for which it still maintains state.

5.1. A Leaf Node Joining a DAG

A new leaf node that joins an established LLN runs an iterative algorithm in which it requests (using multicast DIS) DIOs from routers belonging to the desired DAG. The DIS message has the "No Inconsistency" flag set (to prevent resetting of Trickle timer in responding routers) and the "DIO Type" flag reset (to make responding routers send unicast DIOs back). The DIS message can include a Response Spreading option listing a suitable spreading interval and a Metric Container listing the routing constraints that the responding routers must satisfy. In each iteration, the node multicasts such a DIS and waits for the DIOs. Once the spreading interval has expired, the node considers the current iteration to be unsuccessful. Now the node relaxes the routing constraints somewhat and proceeds to the next iteration. The cycle repeats until the node receives one or more DIOs in a particular iteration or if maximum number of iterations have been reached.

5.2. Identifying A Defunct DAG

An RPL node may remove a neighbor from its parent set for a DAG for a number of reasons:

- o The neighbor is no longer reachable, as determined using a mechanism such as Neighbor Unreachability Detection (NUD) [RFC4861], Bidirectional Forwarding Detection (BFD) [RFC5881] or L2 triggers [RFC5184]; or
- o The neighbor advertises an infinite rank in the DAG; or
- o Keeping the neighbor as a parent would required the node to increase its rank beyond $L + \text{DAGMaxRankIncrease}$, where L is the minimum rank the node has had in this DAG; or
- o The neighbor advertises membership in a different DAG within the same RPL Instance, where a different DAG is recognised by a different DODAGID or a different DODAGVersionNumber.

Even if the conditions listed above exist, an RPL node may fail to remove a neighbor from its parent set because:

- o The node may fail to receive the neighbor's DIOs advertising an increased rank or the neighbor's membership in a different DAG;
- o The node may not check, and hence may not detect, the neighbor's unreachability for a long time. For example, the node may not have any data to send to this neighbor and hence may not encounter any event (such as failure to send data to this neighbor) that would trigger a check for the neighbor's reachability.

In such cases, a node would continue to consider itself attached to a DAG even if all its parents in the DAG are unreachable or have moved to different DAGs. Such a DAG can be characterized as being defunct from the node's perspective. If the node maintains state about a large number of defunct DAGs, such state may prevent a considerable portion of the total memory in the node from being available for more useful purposes.

To alleviate the problem described above, an RPL node may invoke the following procedure to identify a defunct DAG and delete the state it maintains for this DAG. Note that, given the proactive nature of RPL protocol, the lack of data traffic using a DAG can not be considered a reliable indication of the DAG's defunction. Further, the Trickle timer based control of DIO transmissions means the possibility of an indefinite delay in the receipt of a new DIO from a functional DAG parent. Hence, the mechanism described next is based on the use of a DIS message to solicit DIOs about a DAG suspected of defunction. Further, a multicast DIS is used so as to avoid the need to query each parent individually and also to discover other neighbor routers that may serve as the node's new parents in the DAG.

When an RPL node has not received a DIO from any of its parents in a DAG for more than a locally configured time duration:

- o The node generates a multicast DIS message with:
 - * "No Inconsistency" flag set so that the responding routers do not reset their Trickle timers.
 - * "DIO Type" flag set so that the responding routers send multicast DIOs and other nodes in the vicinity do not need to invoke this procedure.
 - * A Solicited Information option to identify the DAG in question. This option must have the I and D flags set and the RPLInstanceID/DODAGID fields must be set to values identifying the DAG. The V flag inside the Solicited Information option should not be set so as to allow the neighbors to send DIOs advertising the latest version of the DAG.

- * A Response Spreading option specifying a suitable time interval over which the DIO responses may arrive.
- o After sending the DIS, the node waits for the duration specified inside the Response Spreading option to receive the DIOs generated by its neighbors. At the conclusion of the wait duration:
 - * If the node has received one or more DIOs advertising newer version(s) of the DAG, it joins the latest version of the DAG, selects a new parent set among the neighbors advertising the latest DAG version and marks the DAG status as functional.
 - * Otherwise, if the node has not received a DIO advertising the current version of the DAG from a neighbor in the parent set, it removes that neighbor from the parent set. As a result, if the node has no parent left in the DAG, it marks the DAG as defunct and schedule the deletion of the state it has maintained for the DAG after a locally configured "hold" duration. (This is because, as per RPL specification, when a node no longer has any parents left in a DAG, it is still required to remember the DAG's identity (RPLInstanceID, DODAGID, DODAGVersionNumber), the lowest rank (L) it has had in this DAG and the DAGMaxRankIncrease value for the DAG for a certain time interval to ensure that the node does not join an earlier version of the DAG and does not rejoin the current version of the DAG at a rank higher than $L + \text{DAGMaxRankIncrease}$.)

6. IANA Considerations

6.1. DIS Flags

IANA is requested to allocate bits 6 and 7 of the DIS Flag Field to become the "No Inconsistency" and "DIO Type" bits, the functionality of which is described in Section 3 of this document.

Value	Meaning	Reference
6	No Inconsistency	This document
7	DIO Type	This document

6.2. RPL Control Message Options

IANA is requested to allocate a new code point in the "RPL Control Message Options" registry for the "Response Spreading" option, the behavior of which is described in Section 4.2.

Value	Meaning	Reference
0x0A	Response Spreading	This document

RPL Control Message Options

7. Security Considerations

TBA

8. References

8.1. Normative References

- [I-D.ietf-roll-rpl] Winter, T., Thubert, P., Brandt, A., Clausen, T., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., and J. Vasseur, "RPL: IPv6 Routing Protocol for Low power and Lossy Networks", draft-ietf-roll-rpl-19 (work in progress), March 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

- [I-D.ietf-roll-terminology] Vasseur, J., "Terminology in Low power And Lossy Networks", draft-ietf-roll-terminology-05 (work in progress), March 2011.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC5184] Teraoka, F., Gogo, K., Mitsuya, K., Shibui, R., and K. Mitani, "Unified Layer 2 (L2) Abstractions for Layer 3 (L3)-Driven Fast Handover", RFC 5184, May 2008.

[RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, June 2010.

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

Authors' Addresses

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

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

Nicolas Dejean
Elster SAS
Espace Concorde, 120 Impasse JB Say
Perols, 34470
France

Email: nicolas.dejean@coronis.com

Dominique Barthel
France Telecom Orange
28 Chemin Du Vieux Chene, BP 98
Meylan, 38243
France

Email: dominique.barthel@orange-ftgroup.com

Emmanuel Baccelli
INRIA

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

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

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

Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: May 1, 2012

M. Goyal, Ed.
University of Wisconsin
Milwaukee
E. Baccelli
INRIA
A. Brandt
Sigma Designs
J. Martocci
Johnson Controls
October 29, 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-02

Abstract

This document specifies a mechanism that enables an RPL router to measure the quality of an existing route towards 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 better 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 May 1, 2012.

Copyright Notice

Copyright (c) 2011 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. Overview	4
3. The Measurement Object (MO)	4
3.1. Format of the base MO	5
3.2. Secure MO	8
4. Originating a Measurement Request	9
4.1. To Measure A Hop-by-hop Route with a Global RPLInstanceID	9
4.2. To Measure A Hop-by-hop Route with a Local RPLInstanceID	9
4.3. To Measure A Source Route	11
5. Processing a Measurement Request at an Intermediate Router . .	12
5.1. Determining Next Hop For An MO Measuring A Source Route .	13
5.2. Determining Next Hop For An MO Measuring A Hop-by-hop Route	13
6. Processing a Measurement Request at the Target	14
7. Processing a Measurement Reply at the Origin	15
8. Security Considerations	15
9. IANA Considerations	16
10. Acknowledgements	16
11. References	17
11.1. Normative References	17
11.2. Informative References	17
Authors' Addresses	17

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, allows 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. Overview

The mechanism described in this document can be used by an origin in an RPL domain to measure the aggregated values of the routing metrics along a P2P route to a target within the same RPL domain. 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. Optionally, the origin may allow an intermediate route to generate the Measurement Reply if it already knows the relevant routing metric values along rest of the route.

3. The Measurement Object (MO)

This document defines two new RPL Control Message types, the Measurement Object (MO), with code 0x06 (to be confirmed by IANA), and the Secure MO, with code 0x86 (to be confirmed by IANA). An MO serves as both Measurement Request and Measurement Reply.

3.1. Format of the base MO

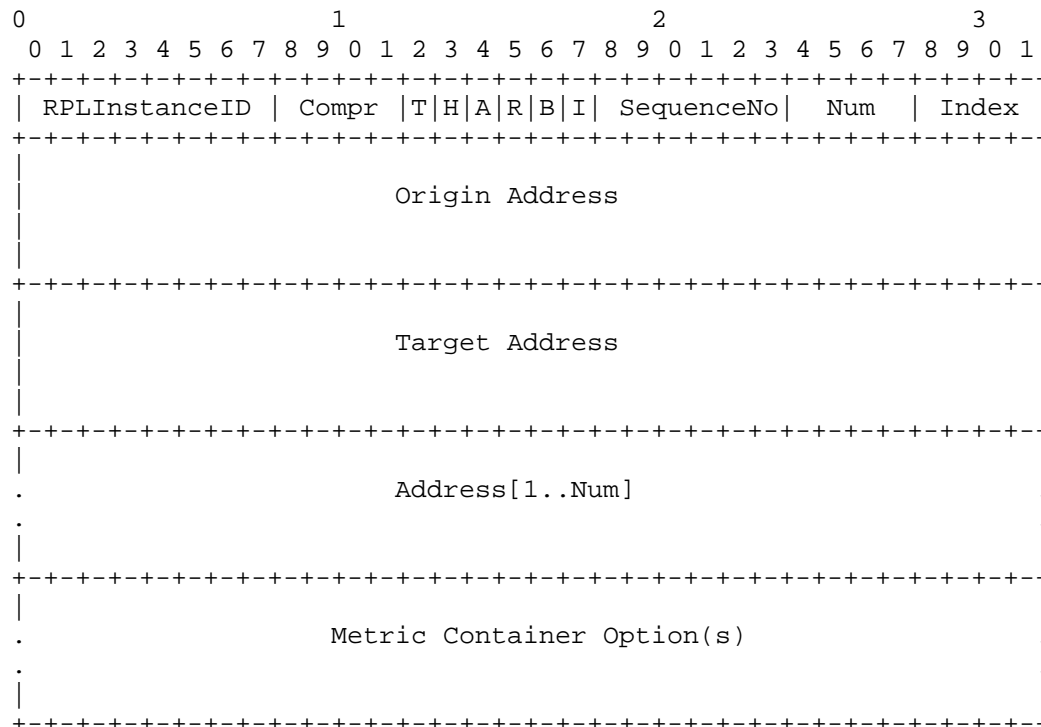


Figure 1: Format of the base Measurement Object (MO)

The format of a base MO is shown in Figure 1. A base 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 being measured. If the route being measured is a source route, this field MUST be set to 10000000 on transmission and ignored on reception.
- o Compr: In many LLN deployments, IPv6 addresses share a well known, common prefix. In such cases, the common prefix can be elided when specifying IPv6 addresses in Origin/Target Address fields and the Address vector. The "Compr" field is a 4-bit unsigned integer that indicates the number of prefix octets that are elided from the IPv6 addresses in Origin/Target Address fields and the Address vector. The Compr value will be 0 if full IPv6 addresses are carried in the Origin/Target Address fields and the Address

vector.

- o Type (T): This flag is set if the MO represents a Measurement Request. The flag is cleared if the MO is a Measurement Reply.
- o 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.
- o 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. In other words, this flag MAY be set only if T = 1, H = 1 and the RPLInstanceID field has a local value. Otherwise, this flag MUST be cleared. A value 1 in this 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.
- o 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. In other words, this flag MAY be set only if T = 1 and H = 0. Otherwise, this flag MUST be cleared. 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.
- o Back Request (B): This flag serves as a request to the target to send a Measurement Request towards the origin. The origin MAY set this flag if it wants to make such a request to the target. On receiving this request, the target MAY generate a Measurement Request to measure the cost of its current (or the most preferred) route to the origin. Receipt of this Measurement Request would allow the origin to know the cost of the back route from the target to itself and thus determine the round-trip cost of reaching the target.

- o Intermediate Reply (I): Relevant only if a hop-by-hop route is being measured, this flag serves as a permission to an intermediate router to generate a Measurement Reply if it knows the cost of the rest of the route being measured. The origin MAY set this flag if a hop-by-hop route is being measured (i.e., H = 1) and the origin wants to allow the intermediate routers to generate the Measurement Reply in response to this Measurement Request. Setting this flag may be useful in scenarios where Hop Count [I-D.ietf-roll-routing-metrics] is the routing metric of interest and the origin expects an intermediate router (e.g. the root of a non-storing DAG or a common ancestor of the origin and the target in a storing DAG) to know the Hop Count of the remainder of the route to the target. This flag MUST be cleared if the route being measured is a source route (i.e., H = 0).
- o SequenceNo: A 6-bit sequence number, assigned by the origin, that allows the origin to uniquely identify a Measurement Request and the corresponding Measurement Reply.
- o 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.
- o Index: If the Measurement Request is traveling along a source route contained in the Address vector (T=1,H=0), 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 (T=1,H=1,A=1 and RPLInstanceID field has a local value), 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 within the RPL domain 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 source route 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 source 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 (origin or an intermediate router) inserting an Address vector inside an MO MUST ensure that no address appears more than once inside the vector. Each router on the way MUST ensure that the loops do not exist within the source route. 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.
- o Metric Container Options: An MO MUST contain one or more Metric Container options to accumulate routing metric values for the route being measured.

3.2. Secure MO

A Secure MO message follows the format in Figure 7 of [I-D.ietf-roll-rpl], where the base format is the base MO shown in Figure 1.

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 below. Additionally, the origin **MUST** set the T flag to 1 to indicate that the MO represents a Measurement Request. The origin **MUST** also 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 below, the origin **MUST** unicast the MO message to the next hop on the P2P route.

4.1. To Measure A Hop-by-hop Route with a Global RPLInstanceID

If a hop-by-hop route with a global RPLInstanceID is being measured, the MO message **MUST NOT** contain the Address vector and the following MO fields **MUST** be set in the manner specified below:

- o Hop-by-hop (H): This flag **MUST** be set;
- o Accumulate Route (A): This flag **MUST** be cleared;
- o Reverse (R): This flag **MUST** be cleared;
- o Back Request (B): This flag **MAY** be set if the origin wants to request the target to generate a Measurement Request back to itself;
- o Intermediate Reply (I): This flag **MAY** be set if the origin wants to permit the intermediate routers to generate the Measurement Reply on the target's behalf;
- o Num: This field **MUST** be set to zero;
- o Index: This field **MUST** be set to zero.

4.2. To Measure A Hop-by-hop Route with a Local RPLInstanceID

If a hop-by-hop route with a local RPLInstanceID is being measured and the MO is not accumulating a source route for the target's use, the MO message **MUST NOT** contain the Address vector and the following MO fields **MUST** be set in the manner specified below:

- o Hop-by-hop (H): This flag **MUST** be set;

- o Accumulate Route (A): This flag MUST be cleared;
- o Reverse (R): This flag MUST be cleared;
- o Back Request (B): This flag MAY be set if the origin wants to request the target to generate a Measurement Request back to itself;
- o Intermediate Reply (I): This flag MAY be set if the origin wants to permit the intermediate routers to generate the Measurement Reply on the target's behalf;
- o Num: This field MUST be set to zero;
- o Index: This field MUST be set to zero;
- o Origin Address: This field MUST contain the DODAGID value (after eliding Compr number of prefix octets) associated with the route being measured.

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 set the following MO fields in the manner specified below:

- o Hop-by-hop (H): This flag MUST be set;
- o Accumulate Route (A): This flag MUST be set;
- o Reverse (R): This flag MUST be cleared;
- o Back Request (B): This flag MAY be set if the origin wants to request the target to generate a Measurement Request back to itself;
- o Intermediate Reply (I): This flag MAY be set if the origin wants to permit the intermediate routers to generate the Measurement Reply on the target's behalf;
- o Address vector: The Address vector must be large enough to accomodate a complete source route from the origin to the target. All the bits in the Address vector field MUST be set to zero;
- o Num: This field MUST specify the number of address elements that can fit inside the Address vector;
- o Index: This field MUST be set to 1;

- o Origin Address: This field MUST contain the DODAGID value (after eliding Compr number of prefix octets) associated with the route being measured.

4.3. To Measure A Source Route

If a source route is being measured, the origin MUST set the following MO fields in the manner specified below:

- o RPLInstanceID: This field MUST be set to 100000000;
- o Hop-by-hop (H): This flag MUST be cleared;
- o Accumulate Route (A): This flag MUST be cleared;
- o Reverse (R): This flag MUST be set if the source route in the Address vector can be reversed and used by the target to source route the Measurement Reply message back to the origin. Otherwise, this flag MUST be cleared;
- o Back Request (B): This flag MAY be set if the origin wants to request the target to generate a Measurement Request back to itself;
- o Intermediate Reply (I): This flag MUST be cleared.
- o Address vector:
 - * The Address vector MUST contain a complete route from the origin to the target (excluding the origin and the target);
 - * The IPv6 addresses (with Compr prefix octets elided) in the Address vector MUST be accessible in the forward direction, i.e., from the origin to the target;
 - * To prevent loops in the source route, the origin MUST ensure that
 - + Any IPv6 address MUST NOT appear more than once in the Address vector;
 - + If the Address vector includes multiple IPv6 addresses assigned to the origin's interfaces, such addresses MUST appear back to back inside the Address vector.
 - * Each address appearing in the Address vector MUST be a unicast address.

- o Num: This field MUST be set to indicate the number of elements in the Address vector;
- o Index: This field MUST be set to 1.

The origin MUST NOT send the packet further if the next hop address on the source route is not on-link.

5. Processing a Measurement Request at an Intermediate Router

A router MAY discard a received MO with no further processing to meet any policy-related goal. Such policy goals may include the need to reduce the router's CPU load or to enhance its battery life.

On receiving an MO, if a router chooses to process the packet further, it MUST check if one of its IPv6 addresses is listed as either the Origin or the Target Address. If not, the router considers itself an Intermediate Router and MUST process the received MO in the following manner.

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

If the I flag is set in the received MO and the intermediate router knows the values of the routing metrics, specified in the Metric Container, for the remainder of the route, it MAY generate a Measurement Reply on the target's behalf in the manner specified in Section 6 (after including in the Measurement Reply the relevant routing metric values for the complete route being measured). Otherwise, the intermediate router MUST process the received MO in the following manner.

The router MUST determine the next hop on the P2P route being measured in the manner described below. The router MUST drop the MO with no further processing and MAY send an ICMPv6 Destination Unreachable (with Code 0 - No Route To Destination) 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 MUST update 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. After updating the routing metrics, the router MUST unicast the MO to the next hop.

5.1. Determining Next Hop For An MO Measuring A Source Route

In case the received MO is measuring a source route ($H=0$), the router MUST increment the Index field and use the Address[Index] element as the next hop. If Index is greater than Num, the router MUST use the Target Address as the next hop.

An intermediate router MUST discard the MO packet with no further processing if the next hop address is not on-link or is not a unicast address. To prevent loops, an intermediate router MUST check if the Address vector includes multiple IPv6 addresses assigned to the router's interfaces and if such addresses do not appear back to back inside the Address vector. In this case, the router MUST discard the MO packet with no further processing. An MO message MUST NOT leave the RPL domain where it originated. Hence, an intermediate router MUST discard an MO message traveling along a source route if the next hop on the way does not lie within the RPL domain.

5.2. Determining Next Hop For An MO Measuring A Hop-by-hop Route

If the received MO is measuring a hop-by-hop route ($H=1$), the router MUST use 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. Moreover,

- o If the RPLInstanceID of the hop-by-hop route is a local value and the A flag is set, the router MUST check if the Address vector already contains one of its IPv6 addresses. If yes, the router MUST discard the packet with no further processing. Otherwise, the router MUST store one of its IPv6 addresses (after eliding Compr prefix octets) at location Address[Index] and then increment 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 as per the following rules:
 - + The Address vector MUST contain a complete route from the router to the target (excluding the router and the target);
 - + The IPv6 addresses (with Compr prefix octets elided) in the Address vector MUST be accessible in the forward direction, i.e., towards the target;

- + To prevent loops in the source route, the router MUST ensure that
 - Any IPv6 address MUST NOT appear more than once in the Address vector;
 - If the Address vector includes multiple IPv6 addresses assigned to the router's interfaces, such addresses MUST appear back to back inside the Address vector.
- + Each address appearing in the Address vector MUST be a unicast address.
- * Specify in the Num field the number of address elements in the Address vector.
- * Set the Index field to 1.

6. Processing a Measurement Request at the Target

On receiving an MO, if a router chooses to process the packet further and finds one of its IPv6 addresses listed as the Target Address, it MUST process the received MO in the following manner.

The target MUST discard the packet with no further processing if the received MO is not a Measurement Request.

The target MUST update the routing metric objects in the Metric Container options if required and MAY note the measured values for the complete route if desired.

The target MUST generate a Measurement Reply message. The received Measurement Request message can be trivially converted into the Measurement Reply by resetting the T flag to zero. The target MAY remove the Address vector from the Measurement Reply if desired. The target MUST then unicast 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 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.

If the B flag is set in the received Measurement Request, the target MAY generate a new Measurement Request to measure the cost of its current (or the most preferred) route to the origin. The routing metrics used in the new Measurement Request MUST include the routing metrics specified in the received Measurement Request.

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 MUST process 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 SHOULD 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 routers on the route, the origin MAY aggregate these local values into an end-to-end value as per the aggregation rules for the metric.

8. Security Considerations

The mechanism defined in this document can potentially be used by a compromised router to generate bogus measurement requests to arbitrary target routers. Such bogus measurement requests may cause processing overload in the routers in the network, drain their batteries and cause traffic congestion in the network. Note that some of these problems would occur even if the compromised router were to generate bogus data traffic to arbitrary destinations.

Since a Measurement Request can travel along a source route specified in the Address vector, some of the security concerns that led to the deprecation of Type 0 routing header [RFC5095] may be valid here. To address such concerns, the mechanism described in this document includes several remedies:

- o This document requires that a route inserted inside the Address vector must be a strict source route and must not include any multicast addresses.
- o This document requires that an MO message must not cross the boundaries of the RPL domain where it is originated. Hence, any security problems associated with the mechanism would be limited to the RPL domain where the MO message is generated.
- o A router must drop a received MO message if the next hop address is not on-link or if it is not a unicast address.
- o A router must check the source route inside the Address vector of each received MO message to ensure that it does not contain a loop involving the router. The router must drop the received packet if the source route does contain such a loop. This and the previous rule protect the network against some of the security concerns even if a compromised node inserts the Address vector inside the MO message.

9. IANA Considerations

IANA is requested to allocate a new code point in the "RPL Control Codes" registry for the "Measurement Object" described in this document.

Code	Description	Reference
0x06	Measurement Object	This document
0x86	Secure Measurement Object	This document

RPL Control Codes

10. Acknowledgements

Authors gratefully acknowledge the contributions of Pascal Thubert, 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., Philipp, M., 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-04 (work in progress), July 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-06 (work in progress), September 2011.
- [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.

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
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

Jerald Martocci
Johnson Controls
507 E Michigan Street
Milwaukee 53202
USA

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

ROLL
Internet-Draft
Intended status: Informational
Expires: April 3, 2012

T. Phinney, Ed.
consultant
P. Thubert
Cisco
RA. Assimiti
Nivis
October 1, 2011

RPL applicability in industrial networks
draft-phinney-roll-rpl-industrial-applicability-00

Abstract

The wide deployment of wireless devices, with their low installed cost (compared to wired devices), will significantly improve the productivity and safety of industrial plants, while simultaneously increasing the efficiency and safety of the plant's workers, by extending and making more timely the information set available about plant operations. The new Routing Protocol for Low Power and Lossy Networks (RPL) defines a Distance Vector protocol that is designed for such networks. The aim of this document is to analyze the applicability of that routing protocol in industrial LLNs of field devices.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <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 3, 2012.

Copyright Notice

Copyright (c) 2011 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	4
2. Terminology	6
3. Overview	7
3.1. Deployment scenarii	7
3.2. Applications and Traffic classes	9
3.3. RPL applicability matrix	10
4. Characterization of communication flows in IACS wireless networks	11
4.1. General	11
4.2. Source-sink (SS) communication paradigm	13
4.3. Publish-subscribe (PS, or pub/sub) communication paradigm	13
4.4. Peer-to-peer (P2P) communication paradigm	15
4.5. Peer-to-multiper (P2MP) communication paradigm	16
4.6. Additional considerations: Duocast and N-cast	17
4.7. RPL applicability per communication paradigm	18
5. RPL profile	21
5.1. Use for process control	21
5.2. RPL features	21
5.2.1. Storing vs. non-storing mode	21
5.2.2. DAO policy	21
5.2.3. Path metrics	22
5.2.4. Objective functions	22
5.2.5. DODAG repair	22
5.2.6. Security	22
5.3. RPL options	22
5.4. Recommended configuration defaults and ranges	22
5.4.1. Trickle parameters	22
5.4.2. Other parameters	23
5.4.3. Additional configuration recommendations	23
6. Other related protocols	24
7. Manageability	25
8. IANA considerations	26
9. Security considerations	27
10. Acknowledgements	28
11. References	29
11.1. Normative References	29
11.2. Informative References	29
11.3. External Informative References	30
Authors' Addresses	31

1. Introduction

Information Technology (IT) is already, and increasingly will be applied to industrial Automation and Control System (IACS) technology in application areas where those IT technologies can be constrained sufficiently by Service Level Agreements (SLA) or other modest change that they are able to meet the operational needs of IACS. When that happens, the IACS benefits from the large intellectual, experiential and training investment that has already occurred in those IT precursors. One can conclude that future reuse of additional IT protocols for IACS will continue to occur due to the significant intellectual, experiential and training economies which result from that reuse.

Following that logic, many vendors are already extending or replacing their local field-bus technology with Ethernet and IP-based solutions. Examples of this evolution include CIP EtherNet/IP, Modbus/TCP, Foundation Fieldbus HSE, PROFINet and Invensys/Foxboro FOXnet. At the same time, wireless, low power field devices are being introduced that facilitate a significant increase in the amount of information which industrial users can collect and the number of control points that can be remotely managed.

IPv6 appears as a core technology at the conjunction of both trends, as illustrated by the current [ISA100.11a] industrial Wireless Sensor Networking (WSN) specification, where layers 1-4 technologies developed for end uses other than IACS - IEEE 802.15.4 PHY and MAC, 6LoWPAN and IPv6, and UDP - are adapted to IACS use. But due to the lack of open standards for routing in Low power and Lossy Networks (LLN), even ISA100.11a leaves the routing operation to proprietary methods.

The IETF ROLL Working Group has defined application-specific routing requirements for a LLN routing protocol, specified in:

Routing Requirements for Urban LLNs [RFC5548],

Industrial Routing Requirements in LLNs [RFC5673],

Home Automation Routing Requirements in LLNs [RFC5826], and

Building Automation Routing Requirements in LLNs [RFC5867].

The Routing Protocol for Low Power and Lossy Networks (RPL) [I-D.ietf-roll-rpl] specification and its point to point extension/optimization [I-D.ietf-roll-p2p-rpl] define a generic Distance Vector protocol that is adapted to a variety of Low Power and Lossy Networks (LLN) types by the application of specific Objective Functions (OFs).

RPL forms Destination Oriented Directed Acyclic Graphs (DODAGs) within instances of the protocol, each instance being associated with an Objective Function to form a routing topology.

A field device that belongs to an instance uses the OF to determine which DODAG and which Version of that DODAG the device should join. The device also uses the OF to select a number of routers within the DODAG current and subsequent Versions to serve as parents or as feasible successors. A new Version of the DODAG is periodically reconstructed to enable a global reoptimization of the graph.

A RPL OF states the outcome of the process used by a RPL node to select and optimize routes within a RPL Instance based on the information objects available. The separation of OFs from the core protocol specification allows RPL to be adapted to meet the different optimization criteria required by the wide range of industrial classes of traffic and applications.

This document provides information on how RPL can accommodate the industrial requirements for LLNs, in particular as specified in [RFC5673].

2. 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 RFC 2119 [RFC2119].

Additionally, this document uses terminology from [I-D.ietf-roll-terminology], and uses usual terminology from the Process Control and Factory Automation industries, some of which is recapitulated below:

FEC: Forward error correction

IACS: Industrial automation and control systems

RAND: reasonable and non-discriminatory (relative to licensing of patents)

3. Overview

3.1. Deployment scenarios

[RFC5673] describes in detail the routing requirements for industrial LLNs. This RFC provides information on the varying deployment scenarios for such LLNs and how RPL assists in meeting those requirements.

Large industrial plants, or major operating areas within such plants, repeatedly go through four major phases, each of which typically lasts from months to years:

P1: Construction or major modification phase

P2: Planned startup phase

P3: Normal operation phase

P4: Planned shutdown phase

followed eventually by an (at least theoretical)

P5: Plant decommissioning phase.

It is also likely, after a major catastrophe at a plant, to have a

P6: Post-emergency recovery and repair phase.

The deployment scenarios for wireless LLN devices may be different in each of these phases. In particular, during the Construction or major modification phase (P1), LLN devices may be installed months before the intended LLN can become usefully operational (because needed routers and infrastructure devices are not yet installed or active), and there are likely to be many personnel in whom the plant owner/operator has only limited trust, such as subcontractors and others in the plant area who have undergone only a cursory background investigation (if any at all). In general, during this phase, plant instrumentation is not yet operational, so could be removed and replaced by a Trojaned device without much likelihood of physical detection of the substitution. Thus physical security of LLN devices is generally a more significant risk factor during this phase than once the plant is operational, where simple replacement of device electronics is detectable.

Extra LLN devices and even extra LLN subnets may be employed during Planned startup (P2) and Planned shutdown (P4) phases, in support of the task of transitioning the plant or plant area between operational

and shutdown states. The extra devices typically provide extra monitoring as the plant transitions infrequent activity states. (In many continuous process plants, up to 2x extra staff are employed at monitoring and control workstations during these two phases, precisely because the plant is undergoing extraordinary behavior as it transitions to or from its steady-state operational condition.)

Similar transient devices and subnets may be used during an unscheduled Post-emergency recovery and repair phase (P6) of operation, but in that case the extra devices usually are routers substituting for plant LLN devices that have been damaged by the incident (such as a fire, explosion, flood, tornado or hurricane) that induced the emergency.

The Planned startup (P2) and Planned shutdown (P4) phases are similar in many respects, but the LLN environment of the two can be quite different, since the Planned shutdown phase can assume that the stable LLN environment used for Normal operation (P3) is functional during shutdown, whereas that stable environment usually is still being established during startup.

The Post-emergency recovery and repair phase (P6) typically operates in an LLN environment that is somewhere between that of the Planned startup (P2) and Normal operation (P3) phases, but with an indeterminate number of temporary routers placed to facilitate communication across and around the area affected by the catastrophe.

Smaller industrial plants and sites may go through similar phases, but often commingle the phases because, in those smaller plants, the phases require less planning and structuring of personnel responsibilities and thus permit less formalization and partitioning of the operating scenarios. For example, it is much simpler, and usually requires much less planning, to bring new equipment on a skid into a plant, using a forklift, than to lay temporary railroad track or employ an extended-axle heavy haul tractor-trailer to deliver a multi-ton process vessel, and temporarily deploy and use very large heavy-lift cranes to install it. In the former cases, nearby equipment usually can continue normal operation while the installation proceeds; in the latter case that is almost always impossible, due to safety and other concerns.

The domain of applicability for the RPL protocol may include all phases but the Normal Operation phase, where the bandwidth allocation and the routes are usually optimized by an external Path Computing Engine (PCE), e.g. an ISA100.11a System Manager.

Additionally, it could be envisioned to include RPL in the normal operation provided that a new Objective Function is defined that

actually interacts with the PCE in order to establish the reference topology, in which case RPL operations would only apply to emergency repair actions. When the reference topology becomes unusable for some failure, and as long as the problem persists.

3.2. Applications and Traffic classes

The industrial market classifies process applications into three broad categories and six classes.

- o Safety
 - * Class 0: Emergency action - Always a critical function
- o Control
 - * Class 1: Closed loop regulatory control - Often a critical function
 - * Class 2: Closed loop supervisory control - Usually non-critical function
 - * Class 3: Open loop control - Operator takes action and controls the actuator (human in the loop)
- o Monitoring
 - * Class 4: Alerting - Short-term operational effect (for example event-based maintenance)
 - * Class 5: Logging and downloading / uploading - No immediate operational consequence (e.g., history collection, sequence-of-events, preventive maintenance)

Safety critical functions effect the basic safety integrity of the plant. These normally dormant functions kick in only when process control systems, or their operators, have failed. By design and by regular interval inspection, they have a well-understood probability of failure on demand in the range of typically once per 10-1000 years.

In-time deliveries of messages becomes more relevant as the class number decreases.

Note that for a control application, the jitter is just as important as latency and has a potential of destabilizing control algorithms.

The domain of applicability for the RPL protocol probably matches the

range of classes where industrial users are interested in deploying wireless networks. This domain includes monitoring classes (4 and 5), and the non-critical portions of control classes (2 and 3). RPL might also be considered as an additional repair mechanism in all situations, and independently of the flow classification and the medium type.

3.3. RPL applicability matrix

It appears from the above sections that whether and the way RPL can be applied for a given flow depends both on the deployment scenario and on the class of application / traffic. At a high level, this can be summarized by the following matrix:

Phase \ Class	0	1	2	3	4	5
Construction			X	X	X	X
Planned startup			X	X	X	X
Normal operation				?	?	?
Planned shutdown			X	X	X	X
Plant decommissioning			X	X	X	X
Recovery and repair	X	X	X	X	X	X

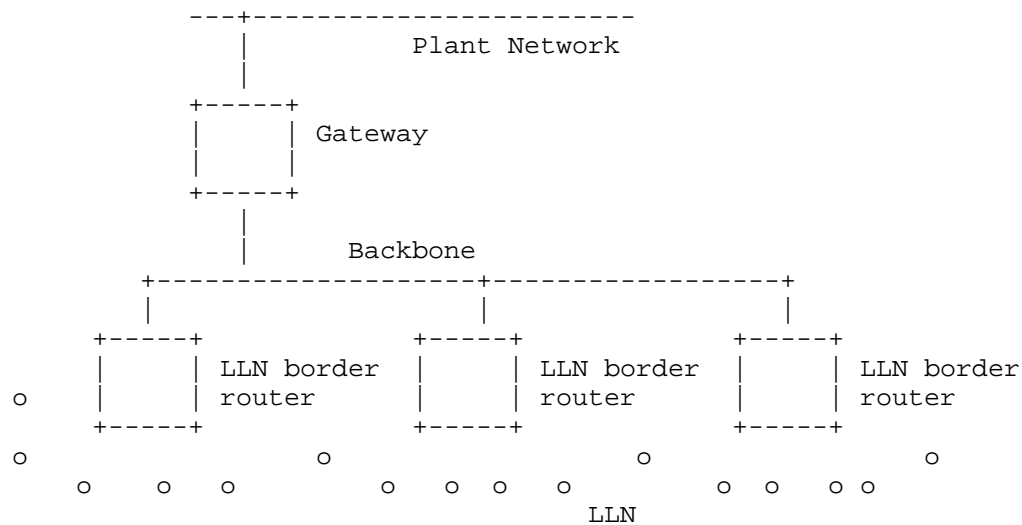
? : typically usable for all but higher-rate classes 0,1 PS traffic

Figure 1: RPL applicability matrix

4. Characterization of communication flows in IACS wireless networks

4.1. General

In an IACS, high-rate communications flows (e.g., 1 Hz or 4 Hz for a traditional process automation network) typically are such that only a single wireless LLN hop separates the source device from a LLN Border Router (LBR) to a significantly higher data-rate backbone network, typically based on IEEE 802.3, IEEE 802.11, or IEEE 802.16, as illustrated in Figure 2.



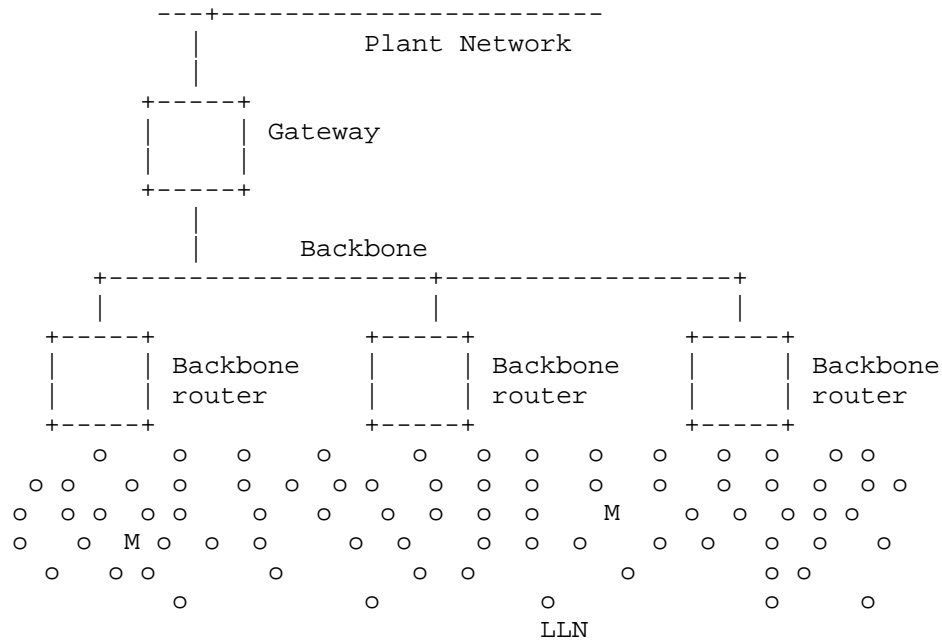
o : stationary wireless field device, seldom acting as an LLN router

Figure 2: High-rate low-delay low-variance IACS topology

For factory automation networks, the basic communications cycle for control is typically much faster, on the order of 100 Hz or more. In this case the LLN itself may be based on high-data-rate IEEE 802.11 or a 100 Mbit/s or faster optical link, and the higher-rate network used by the LBRs to connect the LLN to superior automation equipment typically might be based on fiber-optic IEEE 802.3, with multiple LBRs around the periphery of the factory area, so that most high-rate communications again requires only a single wireless LLN hop.

Multi-hop LLN routing is used within the LLN portion of such networks to provide backup communications paths when primary single-hop LLN paths fail, or for lower repetition rate communications where longer LLN transit times and higher variance are not an issue. Typically,

the majority of devices in an IACS can tolerate such higher-delay higher-variance paths, so routing choices often are driven by energy considerations for the affected devices, rather than simply by IACS performance requirements, as illustrated in Figure 3.



o : stationary wireless field device, often acting as an LLN router
M : mobile wireless device

Figure 3: Low-rate higher-delay higher-variance IACS topology

Two decades of experience with digital fieldbuses has shown that four communications paradigms dominate in IACS:

- SS: Source-sink
- PS: Publish-subscribe
- P2P: Peer-to-peer
- P2MP: Peer-to-multipoint

4.2. Source-sink (SS) communication paradigm

In SS, the source-sink communication paradigm, each of many devices in one set, S1, sends UDP-like messages, usually infrequently and intermittently, to a second set of devices, S2, determined by a common multicast address. A typical example would be that all devices within a given process unit N are configured to send process alarm messages to the multicast address Receivers_of_process_alarms_for_unit_N. Receiving devices, typically on non-LLN networks accessed via LBRs, are configured to receive such multicast messages if their work assignment covers process unit N, and not otherwise.

Timeliness of message delivery is a significant aspect of some SS communication. When the SS traffic conveys process alarms or device alerts, there is often a contractual requirement, and sometimes even a regulatory requirement, on the maximum end-to-end transit delay of the SS message, including both the LLN and non-LLN components of that delay. However, there is no requirement on relative jitter in the delivery of multiple SS messages from the same source, and message reordering during transit is irrelevant.

Within the LLN, the SS paradigm simply requires that messages so addressed be forwarded to the responsible LBR (or set of equivalent LBRs) for further forwarding outside the LLN. Within the LLN such traffic typically is device-to-LBR or device-to-redundant-set-of-equivalent-LBRs. In general, SS traffic may be aggregated before forwarding when both the multicast destination address and other QoS attributes are identical. If information on the target delivery times for SS messages is available to the aggregating forwarding device, that device may intentionally delay forwarding somewhat to facilitate further aggregation, which can significantly reduce LLN alarm-reporting traffic during major plant upset events.

4.3. Publish-subscribe (PS, or pub/sub) communication paradigm

In PS, the publish-subscribe communication paradigm, a device sends UDP-like messages, usually periodically or cyclicly (i.e., repetitively but without fixed periodicity), to a single multicast address derived from or correlated with the device's own address. A typical example would be that each sensor and actuator device within a given process unit N is configured to send process state messages to the multicast address that designates its specific publications. In essence the derived multicast address for device D is Receivers_of_publications_by_device_D. Typically those receivers are in two categories: controllers (C) for control loops in which device D participates, and devices accessed via the LLN's LBRs that monitor and/or accumulate historical information about device D's status and

outputs.

If the controller(s) that receive device D's publication are all outside the LLN and accessed by LBRs, then within the LLN such traffic typically is device-to-LBR or device-to-redundant-set-of-equivalent-LBRs. But if a controller (Cn) is within the LLN, then a number of different LLN-local traffic patterns may be employed, depending on the capabilities of the underlying link technology and on configured performance requirements for such reporting. Typically in such a case, publication by device D is forwarded up a DODAG to an LLN router that is also on a downward DODAG to a destination controller Cn, then forwarded down that second DODAG to that destination controller Cn. Of course, if the LLN router (or even the LBR) is itself the intended destination controller, which will often be the case, then no downward forwarding occurs.

Timeliness of message delivery is a critical aspect of PS communication. Individual messages can be lost without significant impact on the controlled physical process, but typically a sequence of four consecutive lost messages will trigger fallback behavior of the control algorithms, which is considered a system failure by most system owner/operators. (In general, and unless a local catastrophic event such as a major explosion or a tornado occurs in the plant, invocation of more than one instance of such fallback handling per year, per plant, is considered unacceptable.)

Message loss, delay and jitter in delivery of PS messaging is a relative matter. PS messaging is used for transfer of process measurements and associated status from sensors to control computation elements, from control computation elements to actuators, and of current commanded position and status from actuators back to control computation elements. The actual time interval of interest is that which starts with sensing of the physical process (which necessarily occurs before the sensed value can be sent in the first message) and which ends when the computed control correction is applied to the physical process by the appropriate actuator (which cannot occur until after the second message containing the computed control output has been received by that actuator). With rare exception, the control algorithms used with PS messaging in the process automation industries - those managing continuous material flows - rely on fixed-period sampling, computation and transfer of outputs, while those in the factory automation industries - those managing discrete manufacturing operations - rely on bounded delay between sampling of inputs, control computation and transfer of outputs to physical actuators that affect the controlled process.

Deliberately manipulated message delay and jitter in delivery of PS messaging has the potential to destabilize control loops. It is the

responsibility of conveyed higher-level protocols to protect against such potential security attacks by detecting overly delayed or jittered messages at delivery, converting them into instances of message loss. Thus network and data-link protocols such as IPv6 and Ethernet need not themselves address such issues, although their selection and employment should take the existence (or lack) of such higher-layer protection mechanisms, and the resulting consequences due to excessive delay and jitter, into consideration in their parameterization.

In general, PS traffic within the LLN is not aggregated before forwarding, to minimize message loss and delay in reception by any relevant controller(s) that are outside the LLN. However, if all intended destination controllers are within the LLN, and at least one of those intended controllers also serves as an LLN router on a DODAG to off-LLN destinations that all are not controllers, then the router functions in that device may aggregate PS traffic before forwarding when the required routing and other QoS attributes are identical. If information on the target delivery times for PS messages to non-controller devices is available to the aggregating forwarding device, that device may intentionally delay forwarding somewhat to facilitate further aggregation.

In some system architectures, message streams that use PS to convey current process measurements and status are compressed at the source through a 2-dimensional winnowing process that compares

- 1) the process measurement values and status of the about-to-be-sent message with that of the last actually-sent message, and
- 2) the current time vs. the queueing time for the last actually-sent message.

If the interval since that last-sent message is less than a predefined maximum time, and the status is unchanged, and the process measurement(s) conveyed in the message is within predefined deadband(s) of the last-sent measurement value(s), then transmission of the new message is suppressed. Often this suppression takes the form of not queuing the new message for transmission, but in some protocols a brief placeholder message indicating "no significant change" is queued in its stead.

4.4. Peer-to-peer (P2P) communication paradigm

In P2P, the peer-to-peer communication paradigm, a device sends UDP-like or TCP-like messages from one device (D1) to a second device (D2), usually with bidirectional but asymmetric flow of application data, where the amount of data is significantly greater in one

direction than the other. Typical examples are transfer of configuration information to or from a process field device, or transfer of captured process diagnostics (e.g., time-stamped noise signatures from a coriolis flowmeter) to an off-LLN higher-level asset management system. Unicast addressing is used in both directions of data flow.

In general, specific P2P traffic has only loose timeliness requirements, typically just those required so that response times to human-operator-initiated actions meet human factors requirements. As a consequence, in general, message aggregation is permitted, although few opportunities are likely to present themselves for such aggregation due to the sporadic nature of such messaging to a single destination, and/or due to the large message payloads that often occur in at least one direction of transmission.

4.5. Peer-to-multipeer (P2MP) communication paradigm

In P2MP, the peer-to-multipeer communication paradigm, a device sends UDP-like messages downward, from one device (D1) to a set of other devices (Dn). Typical examples are bulk downloads to a set of devices that use identical code image segments or identically-structured database segments; group commands to enable device state transitions that are quasi-synchronized across all or part of the local network (e.g., switch to the next set of point-to-point downloaded session keys, or notifying that the network is switching to an emergency repair and recovery mode); etc. Multicast addressing is used in the downward direction of data flow.

Devices can be assigned to a number of multicast groups, for instance by device type. Then, if it becomes necessary to reflash all devices of a given type with a new load image, a multicast distribution mechanism can be leveraged to optimize the distribution operation.

In general, P2MP traffic has only loose timeliness requirements. As a consequence, in general, message aggregation is permitted, although few opportunities are likely to present themselves for such aggregation due to the sporadic nature of such messaging to a single multicast group destination, and/or due to the large message payloads that often occur when P2MP is used for group downloads. However, in general, message aggregation negatively impacts the delivery success rate for each of the aggregated messages, since the probability of error in a received message increases with message length. Together these considerations often lead to a policy of non-aggregation for P2MP messaging.

Note: Reliable group download protocols, such as the no-longer-published IEEE 802.1E (ISO/IEC 15802-4) system load protocol, and

reliable multicast protocols based on the guidance of RFC2887, are instructive in how P2MP can be used for initial bulk download, followed by either P2MP or P2P selective retransmissions for missed download segments.

4.6. Additional considerations: Duocast and N-cast

In industrial automation systems, some traffic is from (relatively) high-rate monitoring and control loops, of Class 0 and Class 1 as described in [RFC5673]. In such systems, the wireless link protocol, which typically uses immediate in-band acknowledgement to confirm delivery (or, on failure, conclude that a retransmission is required), can be adapted to attempt simultaneous delivery to more than one receiving device, with separated, sequenced immediate in-band acknowledgement by each of those intended receivers. (This mechanism is known colloquially as "duocast" (for two intended receivers), or more generically as "N-cast" (for N intended receivers).) Transmission is deemed successful if at least one such immediate acknowledgement is received by the sending device; otherwise the device queues the message for retransmission, up until the maximum configured number of retries has been attempted.

The logic behind duocast/N-cast is very simple: In wireless systems without FEC (forward error correction), the overall rate of success for transactions consisting of an initial transmission and an immediate acknowledgement is typically 95%. In other words, 5% of such transactions fail, either because the initial message of the transaction is not received correctly by the intended receiver, or because the immediate acknowledgment by that receiver is not received correctly by the transaction initiator.

In the generalized case of N-cast, where any received acknowledgement serves to complete the transaction, and where the N intended receivers are spatially diverse, physically separated from each other by multiple wavelengths, the probability that all such receivers fail to receive the initial message of the transaction, or that all generated immediate acknowledgements are not received by the transaction initiator, is typically approximately $(5\%)^N$. Thus, for duocast, the expected success rate for a single transaction goes from 95% ($1.0 - 0.05$) to 99.75% ($1.0 - 0.05^2$), to 99.9875% ($1.0 - 0.05^3$) when $N=3$, and even higher when $N>3$.

From the above analysis, it is obvious that the primary benefit of N-cast occurs when N goes from $N=1$ (unicast) to $N=2$ (duocast); the reduction in transaction loss rate for increasing $N>2$ is quite small, and for $N>3$ it is infinitesimal. In the typical industrial automation environment of class 1 process control loops, which typically repeat at a 1 Hz or 4 Hz rate, in a very large process

plant with thousands of field devices reporting at that rate, the maximum number of transmission retries that must be planned, and for which capacity must be scheduled (within the requisite 250 ms or 1 s interval) is seven (7) retries for unicast PS reporting, but only three (3) retries with duocast PS reporting. (This is determined by the requirement to not miss four successive reports more than once per year, across the entire plant, as such a loss typically triggers fallback behavior in the controlled loop, which is considered a failure of the wireless system by the plant owner/operator.) In practice, the enormous reduction in both planned and used retransmission capacity provided by duocast/N-cast is what enables 4 Hz loops to be supported in large wireless systems.

When available, duocast/N-cast typically is used only for one-hop PS traffic on Class 1 and Class 0 control loops. It may also be employed for rapid, reliable one-hop delivery of Class 0 and sometimes Class 1 process alarms and device alerts, which use the SS paradigm. Because it requires scheduling of multiple receivers that are prepared to acknowledge the received message during the transaction, in general it is not appropriate for the other types of traffic in such systems - P2P and P2MP - and is not needed for other classes of control loops or other types of traffic, which do not have such stringent reporting requirements.

Note: Although there are known patent applications for duocast and N-cast, at the time of this writing the patent assignee, Honeywell International, has offered to permit cost-free RAND use in those industrial wireless standards that have chosen to employ the technology, under a reciprocal licensing requirement relative to that use. Since duocast and N-cast provide performance and energy optimizations, they are not essential for use in wireless systems. However, in practice, their use makes it possible to support 4 Hz wireless loops and meet sub-second safety alarm reporting requirements in large plants, where that might otherwise be impractical without use of a wired network. When duocast/N-cast is not employed, the wireless retransmission capacity that is needed to support such fast loops often is excessive, typically over 100x that actually used for retransmission (i.e., providing for seven retries per transaction when the mean number used is only 0.06 retries).

4.7. RPL applicability per communication paradigm

To match the requirements above, RPL provides a number of RPL Modes of Operation (MOP):

No downward route: defined in [I-D.ietf-roll-rpl], section 6.3.1, MOP of 0. This mode allows only upward routing, that is from nodes (devices) that reside inside the RPL network toward the outside via the DODAG root.

Non-storing mode: defined in [I-D.ietf-roll-rpl], section 6.3.1, MOP of 1. This mode improves MOP 0 by adding the capability to use source routing from the root towards registered targets within the instance DODAG.

Storing mode without multicast support: defined in [I-D.ietf-roll-rpl], section 6.3.1, MOP of 2. This mode improves MOP 0 by adding the capability to use stateful routing from the root towards registered targets within the instance DODAG.

Storing mode with link-scope multicast DAO: defined in [I-D.ietf-roll-rpl] section 9.10, this mode improves MOP 2 by adding the capability to send Destination Advertisements to all nodes over a single Layer 2 link (e.g. a wireless hop) and enables line-of-sight direct communication.

Storing mode with multicast support: defined in [I-D.ietf-roll-rpl], Mode-of-operation (MOP) of 3. This mode improves MOP 2 by adding the capability to register multicast groups and perform multicast forwarding along the instance DODAG (or a spanning subtree within the DODAG).

Reactive: defined in [I-D.ietf-roll-p2p-rpl], the reactive mode creates on-demand additional DAGs that are used to reach a given node acting as DODAG root within a certain number of hops. This mode can typically be used for an ad-hoc closed-loop communication.

The RPL MOP that can be applied for a given flow depends on the communication paradigm. It must be noted that a DODAG that is used for PS traffic can also be used for SS traffic since the MOP 2 extends the MOP 0, and that a DODAG that is used for P2MP distribution can also be used for downward PS since the MOP 3 extends the MOP 2.

On the other hand, an Objective Function (OF) that optimizes metrics for a pure upwards DODAG might differ from the OF that optimizes a mixed upward and downward DODAG.

As a result, it can be expected that different RPL instances are installed with different OFs, different channel allocations, etc... that result in different routing and forwarding topologies, sometimes with differing delay vs. energy profiles, optimized separately for

the different flows at hand.

This can be broadly summarized in the following table:

Paradigm\RPL MOP	RPL spec	Mode of operation
Peer-to-peer	RPL P2P	reactive (on-demand)
P2P line-of-sight	RPL base	2 (storing) with multicast DAO
P2MP distribution	RPL base	3 (storing with multicast)
Publish-subscribe	RPL base	1 or 2 (storing or not-storing)
Source-sink	RPL base	0 (no downward route)
N-cast publish	RPL base	0 (no downward route)

Figure 4: RPL applicability per communication paradigm

5. RPL profile

5.1. Use for process control

This section outlines a RPL profile for a representative deployment in a process control application. Process monitoring without control is typically less demanding, so a subset of this profile generally will suffice.

5.2. RPL features

5.2.1. Storing vs. non-storing mode

RPL operation is defined for a single RPL instance. However, multiple RPL instances can be supported in multi-service networks where different applications may require the use of different routing metrics and constraints, e.g., a network carrying both safety and non-safety control and monitoring traffic.

In general, storing mode is required for high-reporting-rate devices (where "high rate" is with respect to the underlying link data conveyance capability). Such devices, in the absence of path failure, are typically only one hop from the LBR(s) that convey their messaging to other parts of the system. Fortunately, in such cases, the routing tables required by such nodes are small, even when they include information on DODAGs that are used as backup alternate routes.

In general, devices which communicate with LBRs through a chain of intermediary devices will use storing mode for their upward DODAGs, but will use non-storing mode for downward DODAGs for messaging that they route further into the LLN. However, routers that provide downward forwarding for PS messaging addressed to controllers within the LLN (which is expected to be a rare occurrence) will use storing mode for those forwarding paths, so that timely, destination-constrained forwarding of such recurring messaging does not overload the routing node(s) and their downstream subnets.

5.2.2. DAO policy

Two-way communication is a requirement in industrial automation systems. As a result, nodes SHOULD send DAO messages to establish downward paths from the root to themselves.

<to be added>

5.2.3. Path metrics

RPL relies on an Objective Function for selecting parents and computing path costs and rank. This objective function is decoupled from the core RPL mechanisms and also from the metrics in use in the network. Two objective functions for RPL have been defined at the time of this writing, OF0 and MRHOF, both of which define the selection of a preferred parent and backup parents, and are suitable for industrial automation network deployments.

Neither of the currently defined objective functions supports multiple metrics that might be required in heterogeneous industrial automation networks (e.g., networks composed of devices with different energy and timeliness-of-communication constraints). Additional objective functions specifically designed for such networks may be defined in companion RFCs.

5.2.4. Objective functions

<to be added>

5.2.5. DODAG repair

5.2.6. Security

Industrial automation network deployments typically operate in areas that provide limited physical security (relative to the risk of attack). For this reason, the link layer, transport layer and application layer technologies utilized within such networks typically provide security mechanisms to ensure authentication, confidentiality, integrity, timeliness and freshness. As a result, such deployments may not need to implement RPL's security mechanisms and could rely on link layer and higher layer security features.

5.3. RPL options

5.4. Recommended configuration defaults and ranges

5.4.1. Trickle parameters

Trickle was designed to be density-aware and perform well in networks characterized by a wide range of node densities. The combination of DIO packet suppression and adaptive timers for sending updates allows Trickle to perform well in both sparse and dense environments.

<to be added>

5.4.2. Other parameters

<to be added>

5.4.3. Additional configuration recommendations

<to be added>

6. Other related protocols

<to be added>

7. Manageability

Network manageability is a critical aspect of smart grid network deployment and operation. With millions of devices participating in the smart grid network, many requiring real-time reachability, automatic configuration, and lightweight network health monitoring and management are crucial for achieving network availability and efficient operation.

RPL enables automatic and consistent configuration of RPL routers through parameters specified by the DODAG root and disseminated through DIO packets. The use of Trickle for scheduling DIO transmissions ensures lightweight yet timely propagation of important network and parameter updates and allows network operators to choose the trade-off point they are comfortable with respect to overhead vs. reliability and timeliness of network updates.

The metrics in use in the network along with the Trickle Timer parameters used to control the frequency and redundancy of network updates can be dynamically varied by the root during the lifetime of the network. To that end, all DIO messages SHOULD contain a Metric Container option for disseminating the metrics and metric values used for DODAG setup. In addition, DIO messages SHOULD contain a DODAG Configuration option for disseminating the Trickle Timer parameters throughout the network.

The possibility of dynamically updating the metrics in use in the network as well as the frequency of network updates allows deployment characteristics (e.g., network density) to be discovered during network bring-up and to be used to tailor network parameters once the network is operational rather than having to rely on precise pre-configuration. This also allows the network parameters and the overall routing protocol behavior to evolve during the lifetime of the network.

RPL specifies a number of variables and events that can be tracked for purposes of network fault and performance monitoring of RPL routers. Depending on the memory and processing capabilities of each smart grid device, various subsets of these can be employed in the field.

<to be added>

8. IANA considerations

This specification has no requirement on IANA.

9. Security considerations

This document does not specify operations that could introduce new threats. Security considerations for RPL deployments are to be developed in accordance with recommendations laid out in, for example, [I-D.tsao-roll-security-framework].

Industrial automation networks are subject to stringent security requirements as they are considered a critical infrastructure component. At the same time, since they are composed of large numbers of resource- constrained devices inter-connected with limited-throughput links, many available security mechanisms are not practical for use in such networks. As a result, the choice of security mechanisms is highly dependent on the device and network capabilities characterizing a particular deployment.

In contrast to other types of LLNs, in industrial automation networks centralized administrative control and access to a permanent secure infrastructure is available. As a result link-layer, transport-layer and/or application-layer security mechanisms are typically in place and may make use of RPL's secure mode unnecessary.

10. Acknowledgements

<to be added>

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-of0]
Thubert, P., "RPL Objective Function Zero",
draft-ietf-roll-of0-20 (work in progress), September 2011.
- [I-D.ietf-roll-p2p-rpl]
Goyal, M., Baccelli, E., Philipp, M., 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-04 (work in progress), July 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-06 (work in progress), September 2011.
- [I-D.tsao-roll-security-framework]
Tsao, T., Alexander, R., Daza, V., and A. Lozano, "A Security Framework for Routing over Low Power and Lossy Networks", draft-tsao-roll-security-framework-02 (work in progress), March 2010.
- [RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel, "Routing Requirements for Urban Low-Power and Lossy Networks", RFC 5548, May 2009.
- [RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", RFC 5673, October 2009.
- [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.

11.3. External Informative References

- [HART] www.hartcomm.org, "Highway Addressable Remote Transducer,
a group of specifications for industrial process and
control devices administered by the HART Foundation".

- [ISA100.11a]
ISA, "ISA100, Wireless Systems for Automation", May 2008,
< [http://www.isa.org/Community/
SP100WirelessSystemsforAutomation](http://www.isa.org/Community/SP100WirelessSystemsforAutomation)>.

Authors' Addresses

Tom Phinney (editor)
consultant
5012 W. Torrey Pines Circle
Glendale, AZ 85308-3221
USA

Phone: +1 602 938 3163
Email: tom.phinney@cox.net

Pascal Thubert
Cisco Systems
Village d'Entreprises Green Side
400, Avenue de Roumanille
Batiment T3
Biot - Sophia Antipolis 06410
FRANCE

Phone: +33 497 23 26 34
Email: pthubert@cisco.com

Robert Assimiti
Nivis
1000 Circle 75 Parkway SE, Ste 300
Atlanta, GA 30339
USA

Phone: +1 678 202 6859
Email: robert.assimiti@nivis.com

ROLL
Internet-Draft
Intended status: Standards Track
Expires: April 19, 2012

P. Thubert, Ed.
Cisco Systems
October 17, 2011

RPL adaptation for asymmetrical links
draft-thubert-roll-asymlink-00

Abstract

The Routing Protocol for Low Power and Lossy Networks defines a generic Distance Vector protocol for Low Power and Lossy Networks, many of which exhibit strongly asymmetrical characteristics. This draft proposes an extension for that optimizes RPL operations whereby upwards and downwards direction-optimized RPL instances are associated.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <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 19, 2012.

Copyright Notice

Copyright (c) 2011 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
2. Terminology	3
3. The asymmetrical link problem	4
4. Solution Overview	4
5. Modified DODAG Information Object (DIO)	5
6. Operations	5
7. Backward compatibility	6
8. IANA Considerations	7
9. Security Considerations	7
10. Acknowledgements	7
11. References	8
11.1. Normative References	8
11.2. Informative References	8
Author's Address	8

1. Introduction

The IETF ROLL Working Group has defined application-specific routing requirements for a Low Power and Lossy Network (LLN) routing protocol, specified in [RFC5548], [RFC5673], [RFC5826], and [RFC5867], many of which explicitly or implicitly refer to links with asymmetrical properties.

Upon those requirements, the Routing Protocol for Low Power and Lossy Network [I-D.ietf-roll-rpl] was designed as a platform that can be extended by further specifications or guidances, by adding new metrics, Objective Functions, or additional options.

RPL forms Destination Oriented Directed Acyclic Graphs (DODAGs) within instances of the protocol. Each instance is associated with an Objective Function that is designed to solve the problem that is addressed by that instance.

In one hand, RPL requires bidirectional links for the control, but on the other, there is no requirement that the properties of a link are the same in both directions. In fact, such a symmetry is rarely present in LLNs, whether links are based on radios or power-line.

Some initial implementations require that the quality of both directions of a link is evaluated as very good so that the link can be used for control and data in both directions. This eliminates asymmetrical links that are very good in one direction, but only good enough for scarce activity in the other direction.

In practice, a DAG that is built to optimize upwards traffic is generally not congruent with a DAG that is built to optimize downwards traffic. This is why this specification is designed to enable asymmetrical routing DAGs that are bound together to get the maximum benefits of all bidirectional links.

2. Terminology

The terminology used in this document is consistent with and incorporates that described in 'Terminology in Low power And Lossy Networks' [I-D.ietf-roll-terminology] and [I-D.ietf-roll-rpl].

The term upwards qualifies a link, a DODAG or an instance that is optimal for sending traffic in the general direction of the root, though may be usable but suboptimal for traffic coming from the direction of the root. The term downwards qualifies the same words for the opposite direction.

The term parenting applied to instances refers to the directional association of two instances. The graph formed by parented instances must be a DAG. Traffic may be transferred from an instance onto a parent instance under specified circumstances.

3. The asymmetrical link problem

4. Solution Overview

With the core RPL specification, [I-D.ietf-roll-rpl] each instance is a separate routing topology, and packets must be forwarded within the same topology / same instance. One direct consequence of that design choice is that a topology must be very good for both upwards and downwards traffic; otherwise, traffic between two nodes in the instance may suffer.

A simple approach to address bidirectional but asymmetrical links with RPL is to construct two DAGs, one for upwards traffic and one for downwards traffic, each DAG a separate instance, and then bind the two together. In order to benefit from both instances for a same packet, this solution extends RPL to allow traffic to be transferred from one instance to the next.

It can be noted at this point that with [I-D.ietf-roll-rpl], traffic that goes down does not generally go back up again, whereas P2P traffic within a DODAG might go up to a common parent and then down to the destination. In terms of instance relationship, this means that when an upwards and a downwards instance are bound together, traffic from the former may be transferred to the latter, but not the other way around. In other words, there is an order, a parent-child relationship, between the two instances.

Additionally, if there is no next-hop for a packet going down within the instance, then with [I-D.ietf-roll-rpl] the packet must be dropped. In order to limit that risk, it is tempting though inefficient to lower the constraints that are applied to build the topology. It can be more efficient to actually keep the constraints as they should be, but, instead, enable a less constrained, more spanning, fall-back topology into which traffic can be transferred.

For that reason, this solution allows for more complex instance relationships than plain child-parent associations. In order to avoid loops which could be created when transferring packets from one instance to the next, this solution requires that the instances be themselves organized as a superior Directed Acyclic Graph, and enforce that inter-DAG transfers occur only within that superior

super-DAG of DAG instances.

5. Modified DODAG Information Object (DIO)

The DODAG Information Object [I-D.ietf-roll-rpl] carries information that allows a node to discover a RPL Instance, learn its configuration parameters, select a DODAG parent set, and maintain the DODAG. This specification defines a new flag bit to indicate that the DAG is directional.

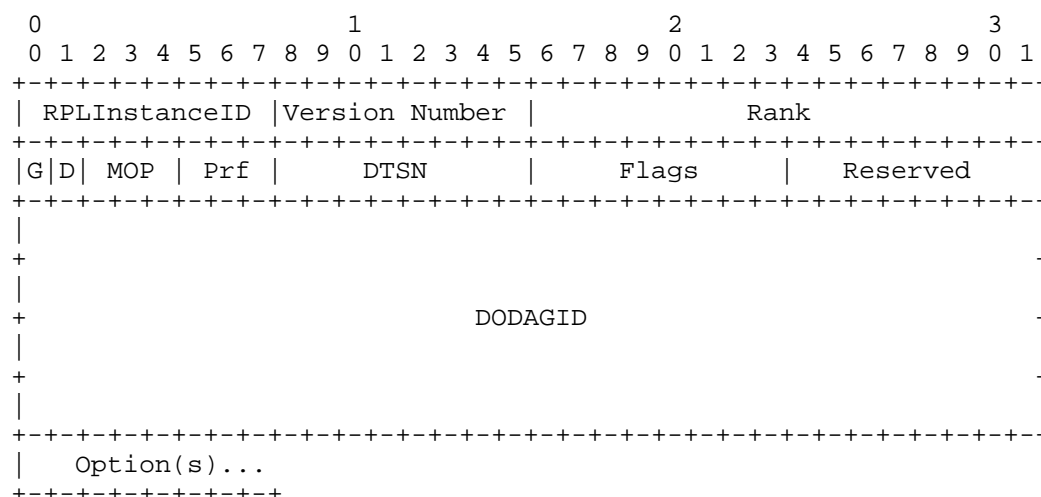


Figure 1: The DIO Base Object

Directional (D): The Directional (D) flag is set to indicate that the instance is intended for directional operation, and reset otherwise. When it is set, a MOP of 0 indicates the upwards direction whereas any other value specified in [I-D.ietf-roll-rpl] indicates downwards. All other values of MOP will be considered downwards unless explicitly specified otherwise.

6. Operations

This specification allows an organization of Instances as follows:

Instances **MUST** be organized as a Directed Acyclic Graph. This information **MUST** be commissioned into the devices so they know both which instances they should participate in, and which

direction of transfer is allowed between instances.

A spanning instance using OF0 [I-D.ietf-roll-of0] MAY be used as root in that instance DAG.

This specification defines a new bit in the RPL [I-D.ietf-roll-rpl] DODAG Information Object (DIO) with the Directional (D) flag that indicates a directional operation for a given instance. An implementation that does not support that new bit will not be able to propagate it.

In case of a directional operation,

The direction is indicated by the MOP field, a MOP of 0 means upwards and otherwise is downwards.

Links are still REQUIRED to allow bidirectional operations

Only the metrics that correspond to the DAG direction are used for the parent selection.

An upward instance SHOULD install routes that lead to the root and beyond - typically the default route.

A downwards instance MAY ONLY install more specific routes that are injected by nodes in the DODAG through the DAO process.

P2P operations are achieved by associating a child upwards instance with a parent downwards instance.

A packet MUST NOT be transferred from a parent instance to a child instance.

A packet MAY be transferred from a child instance to its parent instance if and only if the child instance does not provide a route to the destination, or the parent instance provides a more specific route (longer match) to the destination.

Transferring from an upwards instance to a downwards instance if generally desirable. Other forms of transfers are generally not desirable. Policies MAY be put in place to override that general guidance.

7. Backward compatibility

An OF is generally designed to compute a Rank of a directional link in a fashion that is different from a bidirectional link, and in

particular will not use the same metrics and thus obtain different ranks for a same situation. For that reason, it is important that the OF is aware that an instance is supposed to define a directional DODAG, and it is RECOMMENDED that only devices that support directional DODAGs are allowed in a directional instance.

It might happen that for some purposes like higher availability, an implementation that does not support directional links is administratively allowed to join a directional DODAG. In that case, the extension of the DODAG that starts at that device will not be directional, but the instance will still be functional.

In that case, it might also happen that a device that supports directional DODAGs per this specification sees candidate neighbors that expose the Directional flag and some others that do not. An OF that supports directional links SHOULD favor directional links over non directional links, in a fashion that is to be specified with the OF. In the case of OF0 [I-D.ietf-roll-of0], the 'D' flag should be accounted for before the computation of item 8 in the "Selection Of The Preferred Parent" section 4.2.1., that is before Ranks and be calculated and compared.

8. IANA Considerations

This specification requires that a bit in DIO be assigned to indicate directional link operations as specified in section

9. Security Considerations

Security Considerations for this proposal are to be developed in accordance with recommendations laid out in, for example, [I-D.tsao-roll-security-framework].

10. Acknowledgements

The author wishes to recognize Richard Kelsey, JP Vasseur, Tom Phinney, Robert Assimiti, Don Sturek and Yoav Ben-Yehezkel for their various contributions.

11. References

11.1. Normative References

[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-06 (work in progress), September 2011.

[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-of0]

Thubert, P., "RPL Objective Function Zero", draft-ietf-roll-of0-20 (work in progress), September 2011.

[I-D.tsao-roll-security-framework]

Tsao, T., Alexander, R., Daza, V., and A. Lozano, "A Security Framework for Routing over Low Power and Lossy Networks", draft-tsao-roll-security-framework-02 (work in progress), March 2010.

[RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel, "Routing Requirements for Urban Low-Power and Lossy Networks", RFC 5548, May 2009.

[RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", RFC 5673, October 2009.

[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.

Author's Address

Pascal Thubert (editor)
Cisco Systems
Village d'Entreprises Green Side
400, Avenue de Roumanille
Batiment T3
Biot - Sophia Antipolis 06410
FRANCE

Phone: +33 497 23 26 34
Email: pthubert@cisco.com

