

ROLL
Internet-Draft
Intended status: Standards Track
Expires: September 15, 2011

P. Thubert, Ed.
Cisco Systems
March 14, 2011

RPL Objective Function 0
draft-ietf-roll-of0-07

Abstract

The Routing Protocol for Low Power and Lossy Networks (RPL) defines a generic Distance Vector protocol for Low Power and Lossy Networks (LLNs). RPL is instantiated to honor a particular routing objective/constraint by the adding a specific Objective Function (OF) that is designed to solve that problem. This specification defines a basic OF, OF0, that uses only the abstract properties exposed in RPL messages with no metric container.

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 September 15, 2011.

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. Goal	4
4. Selection of the Preferred Parent	6
5. Selection of the Backup next_hop	7
6. Abstract Interface with RPL core	8
7. OF0 Constants and Variables	8
8. IANA Considerations	9
9. Security Considerations	9
10. Acknowledgements	9
11. References	9
11.1. Normative References	9
11.2. Informative References	9
Author's Address	10

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 [I-D.ietf-roll-building-routing-reqs], [I-D.ietf-roll-home-routing-reqs], [RFC5673], and [RFC5548].

Considering the wide variety of use cases, link types and metrics, the Routing Protocol for Low Power and Lossy Networks [I-D.ietf-roll-rpl] was designed as a generic core that is agnostic to metrics and instantiated using Objective Functions.

RPL forms Destination Oriented Directed Acyclic Graphs (DODAGs) within instances of the protocol, each instance being set up to honor a particular routing objective/constraint of a given deployment. This instantiation is achieved by plugging into the RPL core a specific Objective Function (OF) that is designed to solve that problem to be addressed by that instance.

An Objective Function selects the DODAG version that a device joins, and a number of neighbor routers within that version as parents and siblings. The OF is also responsible for computing the Rank of the device, that abstracts a relative position within the DODAG and is used by the RPL core to enable a degree of loop avoidance and verify forward progression towards a destination, as specified in [I-D.ietf-roll-rpl].

Since there is no default OF or metric container in the RPL main specification, it might happen that, unless given two implementations follow a same guidance for a specific problem or environment, those implementations will not support a common OF with which they could interoperate. This specification fills the need for an Objective Function that can be used as a common denominator between all generic implementations. This is why OF0 is very abstract as to how the link properties are transformed into a Rank, giving only normalized values for what a normal link and what the acceptable range is for a step of Rank are, as opposed to formulating the details of the step of Rank computation.

Indeed, it is the general design in RPL that the metrics are passed from parent to children in a specific container and that the OF will derive the Rank from the natural metric. The separation of Rank and metrics avoids a loss of information as the various metrics are propagated down the DAG. This specification can be used when the link properties that are considered are such that they can be turned in a scalar step of Rank in a reversible fashion and the resulting step of rank is additive over multiple hops.

The Objective Function 0 (OF0) corresponds to the Objective Code Point 0 (OCP0). OF0 does not leverage metric containers such as described in the metrics draft [I-D.ietf-roll-routing-metrics]. OF0 does not require information in the RPL messages but the abstract information from the DIO base container, such as Rank and an administrative preference, that is transported in DIOs as DODAGPreference in [I-D.ietf-roll-rpl]. The Rank of a node is obtained by adding a step of Rank multiplied by a Rank Factor to the Rank of a selected preferred parent. OF0 uses a MinHopRankIncrease of 0x100 so that Rank value can be stored in one octet. This allows up to at least 28 hops even when each hop has the worst step of Rank of 9 and a Rank Factor of 1. How the link properties are transformed into a step of Rank for a given hop depends on the link type and on the implementation. It can be as simple as an administrative cost, but might also derive from a statistical metric with some hysteresis.

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].

3. Goal

The Goal of the OF0 is to join a DODAG version that offers connectivity to a specific set of nodes or to a larger routing infrastructure. For the purpose of OF0, Grounded thus means that the root provides such connectivity. How that connectivity is asserted and maintained is out of scope.

Objective Function 0 is designed to find the nearest Grounded root. In the absence of a Grounded root, LLN inner connectivity is still desirable and floating DAGs will form, rooted at the nodes with the highest administrative preference.

The metric used in OF0 can be an administratively defined scalar cost that is trivially added up along a path to compute the RPL Rank, as defined in [I-D.ietf-roll-rpl]. Depending on how the step of Rank is computed by an implementation, the Rank of a node might be analogous to a weighted hop count of the path to the root. Using a metric that in essence is similar to hop count implies that the quality of the connectivity should be asserted so that only neighbors with a good enough connectivity are presented to the OF. How that connectivity is asserted and maintained is not covered by this specification.

In wireless networks, Hop Count will tend to favor paths with long

distance links and non optimal connectivity properties. In some situations, this might end up partitioning the network. As a result, the link selection must be very conservative, and the available link set is thus constrained. For those reasons, though it can be used on wired links and wired link emulations such as WIFI infrastructure mode, a metric derived from hop count is generally not recommended for wireless networks. Instead, careful thinking should be applied to determine how the step of Rank is computed from the link properties. For instance, the Minimum Rank Objective Function with Hysteresis [I-D.ietf-roll-minrank-hysteresis-of] provides guidance on how hysteresis can be used to maintain a certain stability of the resulting Rank.

The default step of Rank is `DEFAULT_RANK_INCREMENT` for each hop. An implementation MAY allow a step between `MINIMUM_RANK_INCREMENT` and `MAXIMUM_RANK_INCREMENT` to reflect a large variation of link quality by units of `MINIMUM_RANK_INCREMENT`. In other words, the least significant octet in the Rank is not used.

A node MAY stretch its step of Rank by up to `MAXIMUM_RANK_STRETCH` in order to enable the selection of a sibling when only one parent is available. For instance, say that a node computes a step of Rank of 4 units of `MINIMUM_RANK_INCREMENT` from a preferred parent with a Rank of 6 units resulting in a Rank of 10 units for this node. Say that with that Rank of 10 units, this node would end up with only one parent and no sibling, though there is a neighbor with a Rank of 12 units. In that case, the node is entitled to stretch its step of Rank by a value of 2 units, thus using a step of Rank of 6 units so as to reach a Rank of 12 units and find a sibling. But the node is not entitled to use a step of Rank larger than 6 units since that would be a greedy behavior that would deprive the neighbor of this node of a successor. Also, if the neighbor had exposed a Rank of 16 units, the stretch of Rank from 10 to 16 units would have exceeded `MAXIMUM_RANK_STRETCH` of 5 units and thus the neighbor would not have been selectable even as a sibling.

The gap between `MINIMUM_RANK_INCREMENT` and `MAXIMUM_RANK_STRETCH` may not be sufficient in every case to strongly distinguish links of different types or categories in order to favor, say, powered over battery-operated or wired over wireless, within a same DAG. An implementation SHOULD allow a configurable factor called Rank Factor and to apply the factor on all links and peers. An implementation MAY recognize sub-categories of peers and links, such as different MAC types, in which case it SHOULD be able to configure a more specific Rank Factor to those categories. The Rank Factor SHOULD be set between `MINIMUM_RANK_FACTOR` and `MAXIMUM_RANK_FACTOR`. Once a step of Rank is computed along the rules specified in this document, the result of the computation is multiplied by the Rank Factor and the

result is what gets added to the Rank of preferred parent in order to obtain the Rank of this node.

Optionally, the administrative preference of a root MAY be configured to supercede the goal to reach Grounded root. In that case, nodes will associate to the root with the highest preference available, regardless of whether that root is Grounded or not. Compared to a deployment with a multitude of Grounded roots that would result in a same multitude of DODAGs, such a configuration may result in possibly less but larger DODAGs, as many as roots configured with the highest priority in the reachable vicinity.

OF0 selects a preferred parent and a backup next_hop if one is available. The backup next_hop might be but does not have to be a parent or a sibling. All the upward traffic is normally routed via the preferred parent. When the link conditions do not let an upward packet through the preferred parent, the packet is passed to the backup next_hop.

4. Selection of the Preferred Parent

As it scans all the candidate neighbors, OF0 keeps the parent that is the best for the following criteria (in order):

1. [I-D.ietf-roll-rpl] spells out the generic rules for a node to reparent and in particular the boundaries to augment its Rank within a DODAG version. A candidate that would not satisfy those rules MUST NOT be considered.
2. An implementation should validate a router prior to selecting it as preferred. This validation process is implementation and link type dependent, and is out of scope. A router that has been validated is preferable.
3. When multiple interfaces are available, a policy might be locally configured to prioritize them and that policy applies first; that is a router on a higher order interface is preferable.
4. In the absence of a Grounded DODAG version, the router with a higher administrative preference SHOULD be preferred. Optionally, this selection applies regardless of whether the DODAG is Grounded or not.
5. A router that offers connectivity to a grounded DODAG version SHOULD be preferred over one that does not.

6. When comparing 2 routers that belong to the same DODAG, a router that offers connectivity to the freshest sequence SHOULD be preferred.
 7. When computing a resulting Rank for this node from a parent Rank and a Step of Rank from that parent, the parent that causes the lesser resulting Rank SHOULD be preferred.
 8. A DODAG version for which there is an alternate parent SHOULD be preferred. This check is optional. It is performed by computing the backup next_hop while assuming that the router that is currently examined is finally selected as preferred parent.
 9. The DODAG version that was in use already SHOULD be preferred.
 10. The preferred parent that was in use already SHOULD be preferred.
 11. A router that has announced a DIO message more recently SHOULD be preferred.
5. Selection of the Backup next_hop
- o When multiple interfaces are available, a router on a higher order interface is preferable.
 - o The backup next_hop MUST NOT be the preferred parent.
 - o The backup next_hop MUST be either in the same DODAG version as the preferred parent or in an subsequent version. Note that if the backup next_hop is not from the current version then it can not be used as parent.
 - o A Router with a Rank that is higher than the Rank computed for this node out of the preferred parent SHOULD NOT be selected as parent, to avoid greedy behaviors. It MAY still be selected as sibling if no better Back-up next hop is found.
 - o A router with a lesser Rank SHOULD be preferred.
 - o A router that has been validated as usable by an implementation dependant validation process SHOULD be preferred.
 - o The backup next_hop that was in use already SHOULD be preferred.

6. Abstract Interface with RPL core

Objective Function 0 interacts with the core RPL in the following ways:

Processing DIO: This core RPL triggers the OF when a new DIO was received. OF0 analyses the information in the DIO and may select the source as a parent or sibling.

Providing DAG information The OF0 support can be required to provide the DAG information for a given instance to the RPL core. This includes the material that is contained in a DIO base header.

Providing a Parent List The OF0 support can be required to provide the list of the parents for a given instance to the RPL core. This includes the material that is contained in the transit option for that parent.

Trigger The OF0 support may trigger the RPL core to inform it that a change occurred. This can be used to indicate whether the change requires a new DIO to be fired or whether trickle timers need to be reset.

7. OF0 Constants and Variables

OF0 uses the following constants:

MinHopRankIncrease: 256

DEFAULT_RANK_INCREMENT: 3 * MinHopRankIncrease

MINIMUM_RANK_INCREMENT: 1 * MinHopRankIncrease

MAXIMUM_RANK_INCREMENT: 9 * MinHopRankIncrease

MAXIMUM_RANK_STRETCH: 5 * MinHopRankIncrease

DEFAULT_RANK_FACTOR: 1

MINIMUM_RANK_FACTOR: 1

MAXIMUM_RANK_FACTOR: 4

8. IANA Considerations

This specification requires the assignment of an OCP for OF0. The value of 0 is suggested.

9. Security Considerations

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

10. Acknowledgements

Most specific thanks to Philip Levis for his help in finalizing this document, in particular WRT wireless links, to Tim Winter, JP Vasseur, Julien Abeille, Mathilde Durvy, Teco Boot, Navneet Agarwal and Henning Rogge for in-depth review and first hand implementer's feedback.

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-building-routing-reqs]
Martocci, J., Riou, N., Mil, P., and W. Vermeulen,
"Building Automation Routing Requirements in Low Power and Lossy Networks", draft-ietf-roll-building-routing-reqs-07 (work in progress), September 2009.

[I-D.ietf-roll-home-routing-reqs]
Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low Power and Lossy Networks", draft-ietf-roll-home-routing-reqs-08 (work in progress), September 2009.

[I-D.ietf-roll-minrank-hysteresis-of]
Gnawali, O. and P. Levis, "The Minimum Rank Objective Function with Hysteresis", draft-ietf-roll-minrank-hysteresis-of-01 (work in progress), February 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-18 (work in progress), February 2011.

[I-D.ietf-roll-terminology]

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

[I-D.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.

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

