

ROLL
Internet-Draft
Intended status: Standards Track
Expires: 5 November 2026

C. Gundogan
HAW Hamburg
E. Baccelli
INRIA
G. Z. Papadopoulos, Ed.
IMT Atlantique
4 May 2026

RPL DIS Modifications and Applications
draft-ietf-roll-dis-modifications-02

Abstract

This document augments [RFC6550] by defining new DODAG Information Solicitation (DIS) flags and options that enable a RPL node to exert finer control over how neighboring RPL routers respond to its DIO solicitations. In addition, this document describes several use cases that motivate these DIS extensions and illustrate scenarios in which enhanced control of DIO responses improves network efficiency, responsiveness, and robustness.

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 <https://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 5 November 2026.

Copyright Notice

Copyright (c) 2026 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. RFC 6550 refresher	2
1.2. Undesirable effects	4
1.3. Desired improvements	4
2. Terminology	5
3. DIS Base Object flags	5
4. DIS Options	6
4.1. Metric Container	6
4.2. Response Spreading	7
4.3. DIO Option Request	7
5. Full behavior illustration	8
6. Applications	10
6.1. A Leaf Node Joining a DAG	10
6.2. Identifying A Defunct DAG	11
6.3. Adjacencies probing with RPL	13
6.3.1. Deliberations	14
7. IANA Considerations	14
7.1. DIS Flags	14
7.2. RPL Control Message Options	14
8. Security Considerations	14
9. Acknowledgements	15
10. References	15
10.1. Normative References	15
10.2. Informative References	15
Appendix A. Implementation Status	16
Authors' Addresses	16

1. Introduction

This document augments [RFC6550], the RPL routing protocol specification.

1.1. RFC 6550 refresher

Per [RFC6550], a RPL node can send a DODAG Information Solicitation (DIS) message to solicit DODAG Information Object (DIO) messages from neighbor RPL routers.

A DIS can be multicast to all the routers in range or it can be unicast to a specific neighbor router.

A DIS may carry a Solicited Information option that specifies the predicates of the DAG(s) the soliciting node is interested in. In the absence of such Solicited Information option, the soliciting node is deemed interested in receiving DIOs for all the DAGs known by the solicited router(s).

[RFC6550] requires a router to treat the receipt of a multicast DIS as an inconsistency and hence reset its Trickle timers for the matching DAGs. As a result of the general Trickle timer mechanism, future DIOs will be sent at a higher rate. See [RFC6206] for the specification of Trickle timers and the definition of "inconsistency".

[RFC6550] requires a router that receives a unicast DIS to respond by unicasting a DIO for each matching DAG and to not reset the associated Trickle timer. Such a DIO generated in response to a unicast DIS must contain a Configuration option.

This description is summarized in Table 1.

	Unicast DIS	Multicast DIS
No option present	Unicast DIO, do not reset Trickle timer	Do reset Trickle timer
Solicited Information option present, not matching	Do nothing	Do nothing
Solicited Information option present, matching	Unicast DIO, do not reset Trickle timer	Do reset Trickle timer

Table 1: Router behavior on receiving a DIS, as per [RFC6550]

More precisely, Table 1 describes the behavior of routers for each DAG they belong to. In the general case where multiple RPL instances co-exist in a network, routers will maintain a Trickle timer for the one DAG of each RPL instance they belong to, and nodes may send a DIS with multiple Solicited Information options pertaining to different DAGs or instances. In this more general case, routers will respond for each individual DAG/instance they belong to as per Table 1.

1.2. Undesirable effects

As presented in [Sourailidis2020], there are number of undesirable effects linked to the operation of the DIS control message.

Now, consider a 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 will keep transmitting frequent DIOs for a considerable duration until the Trickle timers again reach long intervals. 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 a 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.

Certain scenarios may require a RPL router to solicit a DIO from a parent by using a unicast DIS. The parent is forced to include a Configuration option within the unicast DIO, although the requesting node might still have this information locally available. Since the information within the Configuration option is described as generally static and unchanging throughout the DODAG, it inflates the unicast DIO unnecessarily by 16 bytes for each request.

1.3. Desired improvements

To deal with the situations described above, there is a need in the industry for DIS flags and options that allow a RPL node to control how neighbor RPL routers respond to its solicitation for DIOs, for example by expressing:

- * the routing constraints that routers should meet to be allowed to respond, thereby lowering the number of responders.

- * whether the responding routers should reset their Trickle timers or not, thereby limiting the cumulated number of transmitted DIOs.
- * whether the responding routers should respond with a unicast DIO instead of a multicast one, thereby lowering the overhearing cost in the network.
- * whether the responding routers should omit DIO options that were not requested explicitly and thus reducing the amount of traffic and giving full control over the options of the solicited DIO.
- * the time interval over which the responding routers should schedule their DIO transmissions, thereby lowering the occurrence of collisions.

These results have been attained by the modification implemented and presented in [Sourailidis2020].

2. Terminology

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 [RFC2119].

3. DIS Base Object flags

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

- * the "No Inconsistency" (N) flag: On receiving a multicast DIS with the N flag set, a RPL router MUST NOT reset the Trickle timers for the matching DAGs. In addition, it MUST take specific action, which is to respond by explicitly sending a DIO. This DIO MUST include a Configuration option. This behavior augments [RFC6550], which had provision for such flag. Since this specific, one-shot DIO is not a consequence of the general Trickle timer mechanism, it will be sent right away if no Response Spreading option is present or it will be scheduled according to the Response Spreading option if one is present in the DIS (see Section 4.2).
- * the "DIO Type" (T) flag: In case the N flag is set, this T flag specifies what type of DIO is sent in response. It MUST be a unicast DIO if this flag is set and it MUST be a multicast DIO if this flag is reset.
- * the "DIO Option Request" (R) flag: On receiving a DIS with the R flag set, the receiver MUST include all options that were requested by the DIS containing one or multiple DIO Option Request options. A responding RPL router MUST NOT include DIO options

that were not explicitly requested. Note that this behaviour contradicts with [RFC6550] for the case of including a Configuration option in all DIOs requested by a unicast DIS.

When a unicast DIS is transmitted, both its N and T flags SHOULD be 0, which are the default values per [RFC6550]. On receiving a unicast DIS, the N and T flags MUST be ignored and treated as 00. When the R flag is unset, then a RPL router may include or omit DIO options like specified in [RFC6550]. A RPL router responding to a DIS with the R flag set MUST only include all requested DIO options in the solicited DIO.

The modified DIS base object is shown in Figure 1.




Figure 1: Modified DIS Base Object

4. DIS Options

4.1. Metric Container

In order to lower the number of routers that will respond to a DIS, this document allows routing constraints to be carried by a DIS. Only the router(s) that satisfy these constraints is (are) allowed to respond to the DIS.

These routing constraints are described using a Metric Container option contained in the DIS. Metric Containers are defined in [RFC6550] and [RFC6551]. Metric Containers options were previously only allowed in DIOs. This document augments [RFC6550] by allowing the inclusion of a Metric Container option inside a DIS as well.

A 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 value of the corresponding routing metric that the router maintains for the matching DAG(s). These routing metric values MUST satisfy all the mandatory constraints in order for the router to consider the solicitation successful for the matching DAG(s). This augments the behavior already present in [RFC6550] with the Solicited Information option.

This option can be used in both unicast and multicast DIS.

4.2. Response Spreading




Figure 2: The Response Spreading option

Even with the use of the Solicited Information and the Section 4.1 options, a multicast DIS may still lead to a large number of RPL routers taking immediate action and responding with DIOs. Concurrent transmissions by multiple routers are not desirable since they may lead to poor channel utilization or even to packet loss. Unicast DIOs may be able to avail of link-level retransmissions. However, multicast DIOs usually have no such protection, since they commonly make use of link layer broadcast. 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 option, shown in Figure 2, with a recommended Type value 0x0B (to be confirmed by IANA). A RPL router that explicitly responds with a specific, one-shot DIO to a DIS that includes a Response Spreading option, MUST wait for a time uniformly chosen in the interval $[0..2^{\wedge}\text{SpreadingInterval}]$, 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.

A Response Spreading option MAY be included inside a unicast DIS message, but there is no benefit in doing so.

Multiple Response Spreading options SHOULD NOT be used inside a same DIS message.

This mechanism MUST NOT affect the Trickle timer mechanism.

4.3. DIO Option Request




Figure 3: The DIO Option Request option

If a unicast DIS is used to request a DIO, then [RFC6550] mandates that a Configuration option MUST be included in this DIO. The Configuration option contains generally static information that stays unmodified throughout the DAG. For scenarios where a RPL node is already part of a DAG and hence is holding the information that is propagated with the Configuration option, an inclusion of such leads to an unnecessary inflation of 16 bytes for each solicited DIO.

As per [RFC6550], no process is defined to trigger the inclusion of other DIO options in a solicited DIO.

This document defines a new RPL control message option called DIO Option Request option, shown in Figure 3, with a recommended Type value of 0x0C (to be confirmed by IANA). This new option allows full control over the options of the solicited DIO. The target of a unicast or multicast DIS with the R flag set and with one or more DIO Option Request options included, MUST include these requested options in the solicited DIO. For a DIS with the R flag unset, a RPL router behaves like described in [RFC6550] with regard to DIO options.

5. Full behavior illustration

Table 2 and Figure 5 illustrate the normative behavior described in Section 3 and Section 4.1.

	Unicast DIS	Multicast DIS	Multicast DIS	Multicast DIS
		N=0	N=1, T=0	N=1, T=1
No option present	Unicast DIO, do not reset Trickle timer	Do reset Trickle timer	Multicast DIO, do not reset Trickle timer	Unicast DIO, do not reset Trickle timer
Solicited Information/ Metric Container option present, not matching	Do nothing	Do nothing	Do nothing	Do nothing
Solicited Information/ Metric Container option present, matching	Unicast DIO, do not reset Trickle timer	Do reset Trickle timer	Multicast DIO, do not reset Trickle timer	Unicast DIO, do not reset Trickle timer

Table 2: Router behavior on receiving a DIS, as per [RFC6550]

Notice that Table 2 is indeed identical to Table 1 when Metric Container options are not used in DIS.

For the sake of completeness, let's remind here that a specific, one-shot DIO generated in response to a DIS with the R flag unset MUST contain a Configuration option. If the R flag is set, then this DIO contains only explicitly requested DIO options. This DIO's transmission is delayed according to the Delay Spreading option of the DIS, if one such option is present.

6. Applications

This section details some use cases that require DIS modifications compared to the behaviour currently defined in [RFC6550]. The first use case is that of a new leaf node joining an established DAG in an energy efficient manner. The second use case describes why a node might want to use DIS to identify defunct DAGs for which it still maintains state. The third use case describes the need for adjacency probing and how DIS can be used for that.

6.1. A Leaf Node Joining a DAG

This use case is typically of a smart meter being replaced in the field, while a RPL network is operating and stable. The new smart meter must join the network quickly, without draining the energy of the surrounding nodes, be they battery-operated RPL routers or leaf nodes. In this use case, the issues with the current RPL specification are

- * Just waiting for a gratuitous DIO may take a long time if the Trickle timers have relaxed to the steady state. A technician who has just installed the new meter needs to positively assess that the meter has joined the network before it leaves the premise. It is not economically viable to ask the technician to standby the meter until a gratuitous DIO has arrived, which may take hours.
- * If the meter sends a DIS, it needs to do so using multicast, because it has no knowledge of its surroundings. Sending a multicast DIS is considered an inconsistency by the nearby RPL routers. They will reset their Trickle timer to the shortest period. This will trigger sending a stream of DIOs until the Trickle timers relax again. The DIOs will be sent in multicast, which will trigger energy expenditure at nearby nodes, which had no need for the DIOs.

A proposed solution could be the following. 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, thereby keeping the aggregated number of transmissions low. It also has the "DIO Type" flag set to make responding routers send unicast DIOs back, thereby not triggering full reception in nearby nodes that have state-of-the-art radio receivers with hardware-based address filtering.

The DIS message can include a Response Spreading option prescribing a suitable spreading interval based on the expected density of nearby routers and on the expected Layer 2 technology.

The DIS will likely include a Metric Container listing the routing constraints that the responding routers must satisfy in order to be allowed to respond [RFC6551].

At each iteration, the node multicasts such a DIS and waits for forthcoming DIOs. After a time equal to the spreading interval, the node considers the current iteration to be unsuccessful. The node consequently relaxes the routing constraints somewhat and proceeds to the next iteration.

The cycle repeats until the node receives one or more DIOs or until it has relaxed the constraints to the lowest acceptable values.

This algorithm has been proven in the field to be extremely energy-efficient, especially when routers have a wide communication range.

6.2. Identifying A Defunct DAG

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

- * 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
- * The neighbor advertises an infinite rank in the DAG; or
- * Keeping the neighbor as a parent would require the node to increase its rank beyond $L + \text{DAGMaxRankIncrease}$, where L is the minimum rank the node has had in this DAG; or
- * 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, a RPL node may fail to remove a neighbor from its parent set because:

- * The node may fail to receive the neighbor's DIOs advertising an increased rank or the neighbor's membership in a different DAG;

- * 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, a 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 here 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 a RPL node has not received a DIO from any of its parents in a DAG for more than a locally configured time duration:

- * The node generates a multicast DIS message with:
 - the "No Inconsistency" flag set so that the responding routers do not reset their Trickle timers.
 - the "DIO Type" flag not 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.

- * 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 $L + \text{DAGMaxRankIncrease}$.)

6.3. Adjacencies probing with RPL

RPL avoids periodic hello messaging as compared to other distance vector protocols. It uses trickle timer based mechanism to update configuration parameters. This significantly reduces the RPL control overhead. One of the fallout of this design choice is that, in the absence of regular traffic, the adjacencies could not be tested and repaired if broken.

RPL provides a mechanism in the form of unicast DIS to query a particular node for its DIO. A node receiving a unicast DIS MUST respond with a unicast DIO with Configuration Option. This mechanism could as well be made use of for probing adjacencies and certain implementations such as Contiki uses this. The periodicity of the probing is implementation dependent, but the node is expected to invoke probing only when

- * There is no data traffic based on which the links could be tested.
- * There is no L2 feedback. In some case, L2 might provide periodic beacons at link layer and the absence of beacons could be used for link tests.

6.3.1. Deliberations

- * Should the probing scheme be standardized?
- * In some cases using multicast based probing may prove advantageous. Currently RPL does not have multicast based probing. Multicast DIS/DIO may not be suitable for probing because it could possibly lead to change of states.

7. IANA Considerations

7.1. DIS Flags

IANA is requested to allocate bits 0, 1 and 2 of the DIS Flag Field to become the "No Inconsistency", "DIO Type", and "DIO Option Request" bits, the functionality of which is described in Section 3 of this document.

Value	Meaning	Reference
0	No Inconsistency	This document
1	DIO Type	This document
2	DIO Option Request	This document

Figure 4: DIS Flag Field

7.2. 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 and the "DIO Option Request" option, the behavior of which are described in Section 4.2 and Section 4.3, respectively.

Value	Meaning	Reference
0x0B	Response Spreading	This document
0x0C	DIO Option Request	This document

Figure 5: RPL Control Message Options

8. Security Considerations

TBA

9. Acknowledgements

A lot of text in this document originates from now-expired [I-D.goyal-roll-dis-modifications] co-authored with M. Goyal. The requirements and solutions also draw from now-expired [I-D.dejean-roll-selective-dis] co-authored with N. Dejean. Their contribution is deeply acknowledged.

We also thank (TBA) for their useful feedback and discussion.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/info/rfc6550>>.
- [RFC6551] Vasseur, JP., Ed., Kim, M., Ed., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, DOI 10.17487/RFC6551, March 2012, <<https://www.rfc-editor.org/info/rfc6551>>.

10.2. Informative References

- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [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, DOI 10.17487/RFC5184, May 2008, <<https://www.rfc-editor.org/info/rfc5184>>.

- [RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, DOI 10.17487/RFC5881, June 2010, <<https://www.rfc-editor.org/info/rfc5881>>.
- [RFC6206] Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", RFC 6206, DOI 10.17487/RFC6206, March 2011, <<https://www.rfc-editor.org/info/rfc6206>>.
- [Sourailidis2020]
Sourailidis, D., Koutsiamanis, R., Papadopoulos, G. Z., Barthel, D., and N. Montavont, "RFC 6550: On Minimizing the Control Plane Traffic of RPL-based Industrial Networks", Proceedings of the 21st IEEE International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM), Cork, Ireland, 2020, <<https://doi.org/10.1109/WoWMoM49955.2020.00080>>.

Appendix A. Implementation Status

TBA

Authors' Addresses

Cenk Gundogan
HAW Hamburg
Email: cenk.guendogan@haw-hamburg.de

Emmanuel Baccelli
INRIA
Email: Emmanuel.Baccelli@inria.fr
URI: <https://www.emmanuelbaccelli.org/>

Georgios Z. Papadopoulos (editor)
IMT Atlantique
2 rue de la Chataigneraie
CS 17607
35576 Cesson-Sevigne Cedex
France
Email: georgios.papadopoulos@imt-atlantique.fr

ROLL Working Group
Internet-Draft
Intended status: Standards Track
Expires: 28 July 2026

M. Richardson
Sandelman Software Works
R. A. Jadhav
Huawei Tech
P. Thubert
Cisco Systems
K. Iwanicki
University of Warsaw
24 January 2026

Controlling Secure Network Enrollment in RPL networks
draft-ietf-roll-enrollment-priority-15

Abstract

[RFC9032] defines a method by which a potential [RFC9031] enrollment proxy can announce itself as available for new RPL nodes to enroll on a network. The announcement includes a priority for enrollment. This document provides a mechanism by which a Routing Protocol for Low-Power and Lossy Networks (RPL) Root can globally disable enrollment announcements or adjust the base priority for enrollment operations.

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 <https://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 28 July 2026.

Copyright Notice

Copyright (c) 2026 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Motivation and Overview	2
2. Terminology	3
3. Protocol Definition	4
3.1. Option Format	4
3.2. Option Processing	5
3.3. Upwards Compatibility	6
4. Security Considerations	7
5. Privacy Considerations	8
6. IANA Considerations	8
7. Acknowledgements	8
8. References	8
8.1. Normative References	8
8.2. Informative References	9
Authors' Addresses	9

1. Introduction

[RFC7554] describes the use of the Time-Slotted Channel Hopping (TSCH) mode of [ieee802154]. [RFC9031] and [RFC9032] describe mechanisms by which a new node (the "Pledge") can use a nearby router as a Join Proxy. [RFC9032] describes an extension to the 802.15.4 Enhanced Beacon that is used by a Join Proxy to announce its existence such that Pledges can find them.

1.1. Motivation and Overview

Not every routing member of a mesh ought to announce itself as a `_Join Proxy_`. There are a variety of local reasons for which a 6LowPAN Router (6LR) might not want to provide the `_Join Proxy_` function. Some reasons include low available battery power, already high committed network bandwidth, and lack of available free memory for Neighbor Cache Entry (NCE) slots. An NCE entry is needed in order to maintain communication with the Pledge nodes trying to enroll.

There are other situations where the operator of the network would like to selectively enable or disable the enrollment process in a specific Destination Oriented Directed Acyclic Graph (DODAG). In particular, as the enrollment process involves permitting unencrypted traffic into the best effort part of a network, it would be better to turn the enrollment process off when no new nodes are expected.

This document describes a Routing Protocol for Low-Power and Lossy Networks (RPL) Destination Information Object (DIO) option that can be used to set a minimum enrollment priority. The minimum priority expresses the inability of the RPL DODAG globally to accept new joins. It may derive from multiple constraining factors, for instance, the size of the DODAG, the occupancy of the bandwidth at the DODAG Root, the memory capacity at the Root, or an administrative decision. Each potential `_Join Proxy_` utilizes this value as a base on which to add values relating to local conditions, such as its Rank and number of pending joins. As explained in [RFC9032], higher values decrease the likelihood of an unenrolled node sending enrollment traffic via this `_Join Proxy_`. In particular, by setting the minimum enrollment priority to the maximum value allowed, a network operator can globally disable all new enrollment traffic.

Moreover, when a RPL domain is composed of multiple DODAGs, a node at the edge of more than one such DODAG may not only join any of the DODAGs but also move between them in order to keep their relative sizes balanced. For this, the approximate knowledge of the size of the DODAGs is also an essential metric. Depending on the network policy, the size of the DODAG may or may not affect the minimum enrollment priority. Therefore, since making one proportional to the other would be limiting their value, the current size of the DODAG is advertised separately in the new option.

Updates to the option propagate through the network according to the trickle algorithm. The contents of the option are generated at the DODAG Root and do not change at any hop. If the contents represent an update that is considered important (e.g., quickly disabling any enrollments), the option can trigger trickle timer resets at the nodes to speed up its propagation.

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The term 6LR means 6LowPAN Router, and is defined in [RFC6606]. It refers to a router that forwards packets in a 6LowPAN network.

The terms DAO, DODAG, DODAG root, DIO, trickle timer are from [RFC6550]. The lollipop counter function comes from [RFC6550], Section 7.2.

The term (1)"Join" has been used in documents such as [RFC9031] to denote the activity of a new node authenticating itself to the network to obtain authorization to become a member of the network.

In the context of the [RFC6550] RPL protocol, the term (2)"Join" has an alternative meaning: that of a node (already authenticated to the network, and already authorized to be a member of the network), deciding which part of the RPL DODAG to attach to. This term "Join" has to do with preferred parent selection processes.


In order to avoid the ambiguity of this term, this document refers to the process (1)"Join" as enrollment, leaving the term "Join" to mean (2)"Join". The term "onboarding" (or "IoT Onboarding") is increasingly used to describe what is now called (1)Join in other documents, and is called enrollment in this document. However, the term _Join Proxy_ is retained with its (1)"Join" meaning from [RFC9031].

3. Protocol Definition

This document uses the extensions mechanism specified by [RFC6550]. No mechanism is needed to enable it.

3.1. Option Format

The following option is defined for transmission in DIOs issued by the DODAG Root to be propagated within the DODAG.



Type To be assigned by IANA.

Version Number An 8-bit unsigned integer set by the DODAG root and

denoting the version number of the contents of the option. The version number is interpreted as a lollipop counter (see Section 7.2 of [RFC6550]).

T A bit indicating whether the particular version of the option is important in that adopting its contents should trigger a trickle timer reset at the node.

Min Priority A 7-bit field providing a base value for the Enhanced Beacon Join priority. A value of 0x7f (127) disables the `_Join Proxy_` function entirely.

Exp A 4-bit unsigned integer indicating the power of 2 that defines the unit of the DODAG Size, such that $(\text{unit} = 2^{\text{Exp}})$.

DODAGSz A 4-bit unsigned integer expressing the size of the DODAG in units that depend on the Exp field.

The DODAG Size is calculated as $(\text{DODAGSz} * 2^{\text{Exp}})$.

The DODAG Size can be measured by the Root based on the DAO activity. In such a case, it represents the number of routes not the number of nodes, and can thus be used to infer the load only in a network where each node advertises roughly the same number of addresses and generates roughly the same amount of traffic.

As the DODAG Size is always a multiple of a power of 2, when the actual size falls between two such values, the DODAG Root is to always round up.

Future work such as [I-D.ietf-roll-capabilities] will enable collection of capabilities such as this one in reports to the DODAG Root.

In any case, the DODAG Size may slightly change between one DIO and the next, so the value transmitted is considered as an approximation.

3.2. Option Processing

The contents of the option MUST be generated by the DODAG Root. A 6LR MUST NOT change the option when propagating it.

Whenever the DODAG root changes the values of Min Priority or DODAG Size in the option, it MUST also increment the value of Version Number. Moreover, if the change is considered important (i.e., it is expected to propagate in the DODAG quickly), the DODAG Root SHOULD also set the T bit to 1; otherwise, it MUST set the bit to 0.

Upon receiving the option, a 6LR first checks the value of the Version Number field in the option, `_vr_`, versus the value of the Version Number it has last adopted locally, `_vl_`.

- * If `_vl_` is greater than `_vr_` (in the lollipop counter order), then the 6LR MUST ignore the received option.
- * Otherwise, the 6LR MUST adopt the contents of the option (i.e., the values of Version Number, Min Priority, DODAG Size, and the T bit) as its local ones. Moreover, if `_vl_` was smaller than `_vr_` (in the lollipop counter order) and the T bit in the received option was set, then the 6LR MUST reset its DIO trickle timer.

A 6LR, which would otherwise be willing to act as a `_Join Proxy_`, will examine the locally adopted value of Min Priority and to that number add any additional local consideration (such as upstream congestion, number of NCE slots available, etc.).

The maximum resulting value any 6LR can obtain this way is `0x7f`.

The resulting priority, if less than `0x7f`, should enable the `_Join Proxy_` function.

3.3. Upwards Compatibility

A 6LR that did not support this option would not act on it or propagate it in its DIO messages. In effect, the 6LR's subtree nodes could not receive any telemetry. Therefore, 6LRs that support this option but do not receive it via any path SHOULD assume a default value of `0x40` as their base value for the Enhanced Beacon Join Priority.

A 6LR downstream of a 6LR where there was such an interruption in the telemetry could err in two directions:

- * If the value implied by the base value of `0x40` was too low, then the 6LR might continue to attract enrollment traffic when none should have been collected. This is a stressor for the network, but this would also be what would occur without this option at all.

- * If the value implied by the base value of 0x40 was too high, then the 6LR might deflect enrollment traffic to other parts of the DODAG, possibly refusing any enrollment traffic at all. In order for this to happen, some significant congestion must be seen in the sub-DODAG where the implied 0x40 was introduced. The 0x40 is only the half-way point, so if such an amount of congestion was present, then this sub-DODAG of the DODAG simply winds up being more cautious than it needed to be.

It is possible that the temporal alternation of the above two situations might introduce cycles of accepting and then rejecting enrollment traffic. This is something an operator should consider if they incrementally deploy this option to an existing Low-power/Lossy-Network (LLN). In addition, an operator would be unable to turn off enrollment traffic by sending a maximum value enrollment priority to the sub-DODAG. This situation is unfortunate, but without this option, the situation would occur all over the DODAG, rather than just in the sub-DODAG that the option did not reach.

4. Security Considerations

As per [RFC7416], RPL control frames either run over a secured layer 2 or use the [RFC6550] Secure DIO methods at layer 3. This option can be placed into either a "clear" (layer-2 secured) DIO or a layer-3 Secure DIO.

In most deployments involving wireless technology, layer 2 is always encrypted using a layer-2 specific technology, and so privacy of this option is available.

However, a malicious node that was part of the RPL control plane (i.e., had been enrolled into the layer-2 security) would be able to see the values of this option and, based upon the observed minimal enrollment priority, could signal a confederate that it was a good time to send malicious join traffic.

What is more, such a malicious node, being already part of the RPL control plane, could also send DIOs with a different minimal enrollment priority, which would cause downstream mesh routers to change their `_Join Proxy_` behavior: lower minimal priorities would cause downstream nodes to accept more Pledges than the network was expecting; higher minimal priorities could cause the enrollment process to stall.

The use of layer-2 or layer-3 security for RPL control messages prevents the two aforementioned attacks by non-participating nodes by preventing malicious nodes from becoming part of the control plane.

Nevertheless, a node that is attacked and has malware placed on it creates vulnerabilities in the same way such an attack on any node involved in Internet routing protocol does. The rekeying provisions of [RFC9031] exist to permit an operator to remove such nodes from the network.

5. Privacy Considerations

There are no new privacy issues caused by this extension.

6. IANA Considerations

Allocate a new number TBD01 from Registry RPL Control Message Options. This entry should be called Minimum Enrollment Priority.

7. Acknowledgements

This has been reviewed by Charlie Perkins, Rifaat Shehk-Yusek, Dave Thaler, and Thomas Watteyne.

Huimin She contributed text about expressing the DODAG size.

8. References

8.1. Normative References

- [ieee802154] IEEE standard for Information Technology, "IEEE Std. 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks", n.d., <<http://standards.ieee.org/findstds/standard/802.15.4-2015.html>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/rfc/rfc2119>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/rfc/rfc6550>>.

- [RFC7416] Tsao, T., Alexander, R., Dohler, M., Daza, V., Lozano, A., and M. Richardson, Ed., "A Security Threat Analysis for the Routing Protocol for Low-Power and Lossy Networks (RPLs)", RFC 7416, DOI 10.17487/RFC7416, January 2015, <<https://www.rfc-editor.org/rfc/rfc7416>>.
- [RFC7554] Watteyne, T., Ed., Palattella, M., and L. Grieco, "Using IEEE 802.15.4e Time-Slotted Channel Hopping (TSCH) in the Internet of Things (IoT): Problem Statement", RFC 7554, DOI 10.17487/RFC7554, May 2015, <<https://www.rfc-editor.org/rfc/rfc7554>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.
- [RFC9031] VuiniÄ, M., Ed., Simon, J., Pister, K., and M. Richardson, "Constrained Join Protocol (CoJP) for 6TiSCH", RFC 9031, DOI 10.17487/RFC9031, May 2021, <<https://www.rfc-editor.org/rfc/rfc9031>>.
- [RFC9032] Dujovne, D., Ed. and M. Richardson, "Encapsulation of 6TiSCH Join and Enrollment Information Elements", RFC 9032, DOI 10.17487/RFC9032, May 2021, <<https://www.rfc-editor.org/rfc/rfc9032>>.

8.2. Informative References

- [I-D.ietf-roll-capabilities] Jadhav, R., Thubert, P., Richardson, M., and R. N. Sahoo, "RPL Capabilities", Work in Progress, Internet-Draft, draft-ietf-roll-capabilities-09, 9 November 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-roll-capabilities-09>>.
- [RFC6606] Kim, E., Kaspar, D., Gomez, C., and C. Bormann, "Problem Statement and Requirements for IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing", RFC 6606, DOI 10.17487/RFC6606, May 2012, <<https://www.rfc-editor.org/rfc/rfc6606>>.

Authors' Addresses

Michael Richardson
Sandelman Software Works
Email: mcr+ietf@sandelman.ca

Rahul Arvind Jadhav
Huawei Tech
Email: rahul.ietf@gmail.com

Pascal Thubert
Cisco Systems
Email: pthubert@cisco.com

Konrad Iwanicki
University of Warsaw
Email: iwanicki@mimuw.edu.pl

ROLL Working Group
Internet-Draft
Intended status: Standards Track
Expires: 8 January 2026

R. Koutsiamanis, Ed.
G. Z. Papadopoulos
N. Montavont
IMT Atlantique
P. Thubert
Cisco
7 July 2025

Common Ancestor Objective Function and Parent Set DAG Metric Container
Extension
draft-ietf-roll-nsa-extension-13

Abstract

High reliability and low jitter can be achieved by being able to send data packets through multiple paths, via different parents, in a network. This document details how to exchange the necessary information within RPL control packets to let a node better select the different parents that will be used to forward a packet over different paths. This document also describes the Objective Function which takes advantage of this information to implement multi-path routing.

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 <https://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 8 January 2026.

Copyright Notice

Copyright (c) 2025 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Common Ancestor AP Selection Policies	4
3.1. Common Ancestor Strict	5
3.2. Common Ancestor Medium	6
3.3. Common Ancestor Relaxed	6
4. Common Ancestor Objective Function	6
4.1. Usage	9
5. Node State and Attribute (NSA) object type extension	9
5.1. Usage	11
6. Controlling PRE	12
7. Security Considerations	12
8. IANA Considerations	12
9. Acknowledgments	13
10. References	13
10.1. Normative References	13
10.2. Informative References	13
Appendix A. Implementation Status	14
Appendix B. Choosing an AP selection policy	17
Authors' Addresses	17

1. Introduction

Networks in the industrial context must provide stringent guarantees in terms of reliability and predictability, with this domain being one of the main ones addressed by Deterministic Networking [RFC8557]. One of the ways of achieving such guarantees is through Packet Replication and Elimination (PRE) ([RFC9030], Section 4.5.3), a technique which allows redundant paths in the network to be utilized for traffic requiring higher reliability. Another is to have pre-selected backup paths on standby for quick packet retransmission when packet failures occur. Load-balancing can be also used to make sure that not all traffic passes through the same nodes, to more evenly spread the packet forwarding load. Allowing industrial applications to function over wireless networks requires the application of the principles and architecture of Deterministic Networking [RFC8655]. This results in designs that aim at optimizing packet delivery rate

and bounding latency. Additionally, nodes operating on battery need to minimize their energy consumption.

As an example, to meet this goal, IEEE Std. 802.15.4 [IEEE802154] provides Time-Slotted Channel Hopping (TSCH), a mode of operation that uses a common communication schedule based on timeslots to allow deterministic medium access as well as channel hopping to work around radio interference. However, since TSCH uses retransmissions in the event of a failed transmission, end-to-end latency and jitter performance can deteriorate.

Furthermore, the 6TiSCH working group, focusing on IPv6 over IEEE Std. 802.15.4-TSCH, has worked on these issues and produced the "6TiSCH Architecture" [RFC9030] to address that case.

Building a multi-path DODAG can be achieved based on the RPL capability of having multiple parents for each node in a network, a subset of which is used to forward packets. In order to select parents to be part of this subset, the RPL Objective Function (OF) needs additional information. This document describes an OF which implements multi-path routing and specifies the transmission of this specific path information.

This document describes a new Objective Function (OF) called the Common Ancestor (CA) OF (see Section 4). A detailed description is given of how the path information is used within the CA OF and how the subset of parents for forwarding packets is selected. This specification defines a new Objective Code Point (OCP) for the CA OF.

For the path information, this specification focuses on the extensions to the DAG Metric Container [RFC6551] required for supplying to the CA OF a part of the information it needs to operate. This information is the RPL [RFC6550] parent address set of a node and it must be sent to potential children of the node. The RPL DIO Control Message is the canonical way of broadcasting this kind of information and therefore its DAG Metric Container [RFC6551] field is used to append a Node State and Attribute (NSA) object. The node's parent address set is stored as an optional TLV within the NSA object. This specification defines the type value and structure for the parent address set TLV (see Section 5).

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

The draft uses the following Terminology from other RFCs:

Parent Set (PS): Defined in RPL [RFC6550].

Packet Replication and Elimination (PRE): A method that consists of transmitting multiple copies of a packet using multi-path forwarding over a multi-hop network and that consolidates multiple received packet copies to control flooding. See "Complex Track with Replication and Elimination" in [RFC9030], Section 4.5.3 for more details.

The draft introduces the following Terminology:

Alternative Parent (AP): An RPL parent in the parent set of a node is used to forward a packet copy when replicating packets.

Alternative Parent (AP) Selection: The mechanism for choosing the next hop node to forward a packet copy when replicating packets.

Preferred Grand Parent (PGP): The preferred parent of the preferred parent of a node.

3. Common Ancestor AP Selection Policies

In the RPL protocol, each node maintains a list of potential parents. When more than one parent is required, as when performing PRE, the RPL DODAG Preferred Parent node is used, as per RPL [RFC6550] parent selection, effectively depending on the OF used. If the CA OF is used, the way this choice is made is described in Section 4. Furthermore, to construct an alternative path toward the root, in addition to the PP node, each node in the network selects one or more parents, called Alternative Parents (APs), from its Parent Set (PS).

There are multiple possible policies for selecting the AP node. This section details three such possible policies.

All three policies defined perform AP selection based on common ancestors, named Common Ancestor Strict, Common Ancestor Medium, and Common Ancestor Relaxed, depending on how restrictive the selection process is. A more restrictive policy will limit flooding but might fail to select an appropriate AP, while a less restrictive one will more often find an appropriate AP but might increase flooding.

All three policies apply their corresponding common ancestor criterion to filter the list of candidate neighbors in the Alternative Parent set.

If after the filtering there are multiple condition-meeting candidate nodes, the node MUST select at least one of them as its AP node. The way this choice is made depends on which OF is used. If the CA OF is used, the way this choice is made is described in Section 4.

3.1. Common Ancestor Strict

In the CA Strict OF the node will check if its Preferred Grand Parent (PGP), the PP of its PP, is the same as the PP of the potential AP.

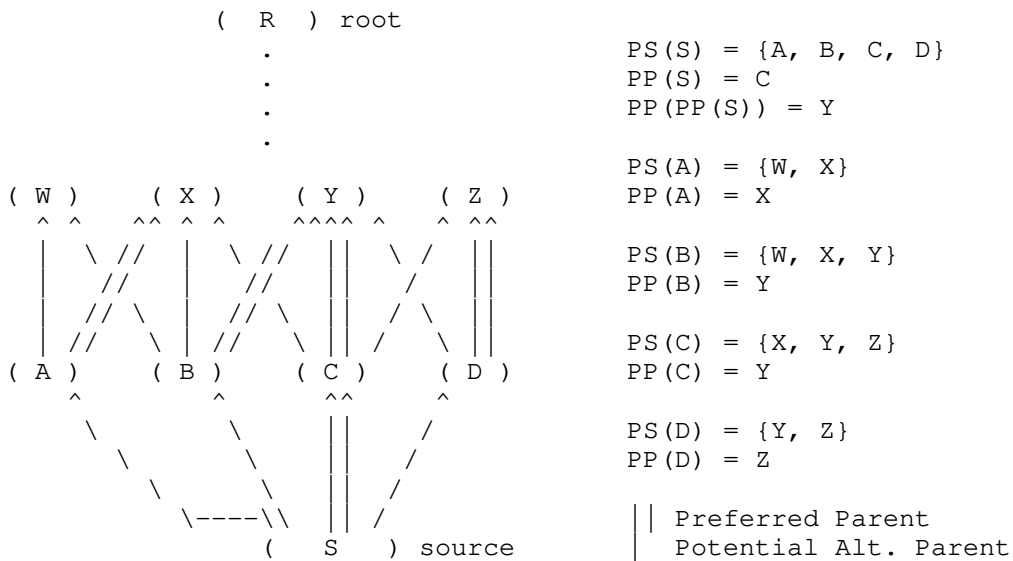


Figure 1: Example Common Ancestor Strict Alternative Parent Selection policy

For example, in Figure 1, the source node S must know its grandparent sets through nodes A, B, C, and D. The Parent Sets (PS) and the Preferred Parents (PP) of nodes A, B, C, and D are shown on the side of the figure. The CA Strict parent selection policy will select an AP for node S for which $PP(PP(S)) = PP(AP)$. Given that $PP(PP(S)) = Y$:

- * Node A: $PP(A) = X$ and therefore it is different than $PP(PP(S))$
- * Node B: $PS(B) = Y$ and therefore it is equal to $PP(PP(S))$
- * Node D: $PS(D) = Z$ and therefore it is different than $PP(PP(S))$

Therefore, node S MUST select node B as its AP node, since $PP(PP(S)) = Y = PP(B)$.

3.2. Common Ancestor Medium

In the CA Medium OF the node will check if its Preferred Grand Parent (PGP), the PP of its PP, is contained in the PS of the potential AP.

Using the same example, in Figure 1, the CA Medium parent selection policy will select an AP for node S for which $PP(PP(S))$ is in $PS(AP)$. Given that $PP(PP(S)) = Y$:

- * Node A: $PS(A) = \{W, X\}$ and therefore $PP(PP(S))$ is not in the set
- * Node B: $PS(B) = \{W, X, Y\}$ and therefore $PP(PP(S))$ is in the set
- * Node D: $PS(D) = \{Y, Z\}$ and therefore $PP(PP(S))$ is in the set

Therefore, S MUST select at least one node among B and D as its AP node.

3.3. Common Ancestor Relaxed

In the CA Relaxed OF the node will check if the Parent Set (PS) of its Preferred Parent (PP) has a node in common with the PS of the potential AP.

Using the same example, in Figure 1, the CA Relaxed parent selection policy will select an AP for node S for which $PS(PP(S))$ has at least one node in common with $PS(AP)$. Given that $PS(PP(S)) = \{X, Y, Z\}$:

- * Node A: $PS(A) = \{W, X\}$ and the common nodes are $\{X\}$
- * Node B: $PS(B) = \{W, X, Y\}$ and the common nodes are $\{X, Y\}$
- * Node D: $PS(D) = \{Y, Z\}$ and the common nodes are $\{Y, Z\}$

Therefore, S MUST select at least one node among A, B, and D as its AP node.

4. Common Ancestor Objective Function

An OF which allows the multiple paths to remain correlated is detailed here. More specifically, when using this OF a node will select an AP node "close" to its PP node to allow the operation of overhearing between parents. Closeness here is not strictly defined, however, the premise is that those candidate parent nodes that have common parents themselves have a higher probability of being within each other's radio range, though it's of course not guaranteed. For more details about overhearing and its use in this context see the "Complex Track with Replication and Elimination" in [RFC9030],

Section 4.5.3. If multiple potential APs match this condition, one of the APs with the lowest rank will be registered, with the choice between multiple nodes with the same lowest rank being implementation-specific.

The OF described here is an extension of The Minimum Rank with Hysteresis Objective Function (MRHOF) [RFC6719]. The CA OF does not update [RFC6719]. Rather, it uses the existing definition of MRHOF in [RFC6719] to build a new OF (with a new Objective Code Point (OCP)) which provides additional functionality, while maintaining compatibility by retaining the existing functionality of MRHOF for the preferred parent. To be precise, this OF extends MRHOF by specifying how an AP is selected while the selection and switching of the PP remain unaltered. Importantly, the calculation of the rank of the node through each candidate neighbor and the selection of the PP is kept the same as in MRHOF.

How the CA OF differs from MRHOF in a section-by-section manner follows in detail:

[RFC6719], Section 2: "Terminology". Term "Selected Metric":

The CA OF uses only one metric, like MRHOF, for rank calculation, with the same MRHOF semantics. For selecting the AP, the PS TLV (stored in the DIO Metric Container Node State and Attribute (NSA) object body, see Section 5) is used. This additional NSA metric is disregarded for rank calculation.

[RFC6719], Section 3 "The Minimum Rank with Hysteresis Objective Function":

Same as MRHOF extended to AP selection. Minimum Rank path selection and switching apply correspondingly to the AP with the extra CA requirement of having some match between ancestors, according to one of the Common Ancestor AP selection policies defined in Section 3.

[RFC6719], Section 3.1 "Computing the Path Cost":

Same as MRHOF extended to AP selection. If a candidate neighbor does not fulfill the CA requirement then the path cost through that neighbor MUST be set to MAX_PATH_COST, the same value used by MRHOF. As a result, the node MUST NOT select the candidate neighbor as its AP.

[RFC6719], Section 3.2 "Parent Selection":

Same as MRHOF extended to AP selection. To allow hysteresis, AP selection maintains a variable, `cur_ap_min_path_cost`, which is the path cost of the current AP.

[RFC6719], Section 3.2.1 "When Parent Selection Runs":

Same as MRHOF.

[RFC6719], Section 3.2.2 "Parent Selection Algorithm":

Same as MRHOF extended to AP selection. If the smallest path cost for paths through the candidate neighbors is smaller than `cur_ap_min_path_cost` by less than `PARENT_SWITCH_THRESHOLD` (the same variable as MRHOF uses), the node MAY continue to use the current AP. Additionally, if there is no PP selected, there MUST NOT be any AP selected either. Finally, as with MRHOF, a node MAY include up to `PARENT_SET_SIZE-1` additional candidate neighbors in its Alternative Parent set. The value of `PARENT_SET_SIZE` is the same as in MRHOF.

[RFC6719], Section 3.3 "Computing Rank":

Same as MRHOF.

[RFC6719], Section 3.4 "Advertising the Path Cost":

Same as MRHOF.

[RFC6719], Section 3.5 "Working without Metric Containers":

The CA OF can work without metric containers identically to MRHOF. Nodes that transmit DIO messages without the Metric Container will never be selected as an AP by the CA OF of another node but can be selected as the PP as per the operation of MRHOF. Effectively, the lack of Metric Containers is equivalent to operating with a Parent Set TLV where there are no PS IPv6 addresses and the PS Length is 0.

[RFC6719], Section 4 "Using MRHOF for Metric Maximization":

Same as MRHOF.

[RFC6719], Section 5 "MRHOF Variables and Parameters":

Same as MRHOF extended to AP selection. The CA OF operates like MRHOF for AP selection by maintaining separate:

AP: Corresponding to the MRHOF PP. Hysteresis is configured for AP with the same `PARENT_SWITCH_THRESHOLD` parameter as in MRHOF. The AP MUST NOT be the same as the PP.

Alternative parent set: Corresponding to the MRHOF parent set. The size is defined by the same `PARENT_SET_SIZE` parameter as in MRHOF. The Alternative parent set MUST be a strict subset of the parent set.

`cur_ap_min_path_cost`: Corresponding to the MRHOF `cur_min_path_cost` variable. To support the operation of the hysteresis function for AP selection.

[RFC6719], Section 6 "Manageability":
Same as MRHOF.

[RFC6719], Section 6.1 "Device Configuration":
Same as MRHOF.

[RFC6719], Section 6.2 "Device Monitoring":
Same as MRHOF.

4.1. Usage

All the Common Ancestor AP Selection Policies (Section 3) apply their corresponding criterion to filter the list of candidate neighbors in the Alternative Parent set. The AP is then selected from the Alternative Parent set based on Rank and using hysteresis as is done for the PP in MRHOF. It is noteworthy that the OF uses the same Objective Code Point (OCP): (TBD1) for all policies used.

The PS information can be used by any of the described AP selection policies or other ones not described here, depending on requirements. It is optional for all nodes to use the same AP selection policies. Different nodes may use different AP selection policies since the selection policy is local to each node. For example, using different policies can be used to vary the transmission reliability in each hop. Some suggestions are provided in Appendix B.

5. Node State and Attribute (NSA) object type extension

In order to select their AP node, nodes need to be aware of their grandparent node sets. Within RPL [RFC6550], the nodes use the DODAG Information Object (DIO) Control Message to broadcast information about themselves to potential children. However, RPL [RFC6550], does not define how to propagate information related to the parent set, which is what this document addresses.

DIO messages can carry multiple options, out of which the DAG Metric Container option [RFC6551] is the most suitable structurally and semantically to carry the parent set. The DAG Metric Container option itself can carry different nested objects, out of which the Node State and Attribute (NSA) [RFC6551] is appropriate for transferring generic node state data. Within the Node State and Attribute, it is possible to store optional TLVs representing various node characteristics. As per the Node State and Attribute (NSA) [RFC6551] description, no TLV has been defined for use. This document defines one TLV for transmitting a node's parent set.

The structure of the DAG Metric Container data in the form of a Node State and Attribute (NSA) object with a TLV in the NSA Optional TLVs field is shown in Figure 3. The first 32 bits comprise the DAG Metric Container header and all the following bits are part of the Node State and Attribute object body, as defined in [RFC6551]. This document defines a new TLV, which MUST be carried in the Node State and Attribute (NSA) object Optional TLVs field within the context of the use of the CA OF. The TLV is named Parent Set and is abbreviated as PS in Figure 3.

PS type: The type of the Parent Set TLV. The value is (TBD2).

PS Length: The total length of the TLV value field (PS IPv6 address(es)) in bytes (0 included). The length is an integral multiple of 16, the number of bytes in an IPv6 address.

PS IPv6 address(es): One or more 128-bit IPv6 addresses, without any separator between them. The field consists of one IPv6 address per parent in the parent set. The parent addresses are listed in decreasing order of preference and not all parents in the parent set need to be included. The selection of how many parents from the parent set will be included is left to the implementation. The number of parent addresses in the PS IPv6 address(es) field can be deduced by dividing the length of the PS IPv6 address(es) field in bytes by 16, the number of bytes in an IPv6 address.

5.1. Usage

The PS is used in the process of parent selection, and especially in AP selection since it can help the alternative path to not significantly deviate from the preferred path. The Parent Set is information local to the node that broadcasts it.

The PS is used only within NSA objects configured as a metric, therefore the DAG Metric Container field "C" MUST be 0. Additionally, since the information in the PS needs to be propagated downstream but cannot be aggregated, the DAG Metric Container field "R" MUST be 1. Finally, since the information contained is by definition partial, specifically just the parent set of the DIO-sending node, the DAG Metric Container field "P" MUST be 1.

The presence of incorrectly configured flags MUST render the Parent Set TLV invalid. This case MUST be handled equivalently to operating with a Parent Set TLV where there are no PS IPv6 addresses and the PS Length is 0.

The presence of a PS Length value that is not a multiple of 16 or larger than 240 MUST render the Parent Set TLV invalid. This case MUST be handled equivalently to operating with a Parent Set TLV where there are no PS IPv6 addresses and the PS Length is 0.

6. Controlling PRE

PRE is very helpful when the aim is to increase reliability for a certain path, however, its use creates additional traffic as part of the replication process. It is conceivable that not all paths have stringent reliability requirements. Therefore, a way to control whether PRE is applied to a path's packets SHOULD be implemented. For example, a traffic class label can be used to determine this behavior per flow type as described in Deterministic Networking Architecture [RFC8655].

7. Security Considerations

All the security considerations from [RFC6550], [RFC6551], and [RFC6719] apply.

In this document, the structure of the DIO control message is extended, within the pre-defined DIO options. The additional information is the list of IPv6 addresses of the parent set of the node transmitting the DIO. This use of this additional information can have the following additional potential consequences:

- * A malicious node that can send DIOs can use the parent set extension to convince neighbors to route through itself, instead of the normal preferred parent they would use. However, this is already possible with other OFs (like OF0 [RFC6552] and MRHOF [RFC6719]) by reporting a fake rank value in the DIO, thus masquerading as the DODAG root.

8. IANA Considerations

This document requests the allocation of a new value (TBD1) from the "Objective Code Point (OCP)" registry in the "Routing Protocol for Low Power and Lossy Networks (RPL)" registry group. The Description field should have the value "Common Ancestor Objective Function (CAOF)".

This document also requests the allocation of a new value (TBD2) for the "Parent Set" TLV from the "Routing Metric/Constraint TLVs" registry in the "Routing Protocol for Low Power and Lossy Networks (RPL) Routing Metric/Constraint" registry group. The Description field should have the value "Parent Set".

9. Acknowledgments

We are very grateful to Dominique Barthel, Rahul Jadhav, Fabrice Theoleyre, Diego Dujovne, Derek Jianqiang Hou, Michael Richardson, and Alvaro Retana for their comments, feedback, and support which lead to many improvements to this document. We would also like to thank Tomas Lagos Jenschke very much for helping in the implementation and evaluation of this document.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/rfc/rfc2119>>.
- [RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <<https://www.rfc-editor.org/rfc/rfc6550>>.
- [RFC6551] Vasseur, JP., Ed., Kim, M., Ed., Pister, K., Dejean, N., and D. Barthel, "Routing Metrics Used for Path Calculation in Low-Power and Lossy Networks", RFC 6551, DOI 10.17487/RFC6551, March 2012, <<https://www.rfc-editor.org/rfc/rfc6551>>.
- [RFC6719] Gnawali, O. and P. Levis, "The Minimum Rank with Hysteresis Objective Function", RFC 6719, DOI 10.17487/RFC6719, September 2012, <<https://www.rfc-editor.org/rfc/rfc6719>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.

10.2. Informative References

- [IEEE802154] IEEE standard for Information Technology, "IEEE Std. 802.15.4, Part. 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks", n.d., <<http://standards.ieee.org/findstds/standard/802.15.4-2015.html>>.
- [RFC6552] Thubert, P., Ed., "Objective Function Zero for the Routing Protocol for Low-Power and Lossy Networks (RPL)", RFC 6552, DOI 10.17487/RFC6552, March 2012, <<https://www.rfc-editor.org/rfc/rfc6552>>.
- [RFC8557] Finn, N. and P. Thubert, "Deterministic Networking Problem Statement", RFC 8557, DOI 10.17487/RFC8557, May 2019, <<https://www.rfc-editor.org/rfc/rfc8557>>.
- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/rfc/rfc8655>>.
- [RFC9030] Thubert, P., Ed., "An Architecture for IPv6 over the Time-Slotted Channel Hopping Mode of IEEE 802.15.4 (6TiSCH)", RFC 9030, DOI 10.17487/RFC9030, May 2021, <<https://www.rfc-editor.org/rfc/rfc9030>>.

Appendix A. Implementation Status

A research-stage implementation of the PRE mechanism using the proposed extension as part of a 6TiSCH IOT use case was developed at IMT Atlantique, France by Tomas Lagos Jenschke and Remous-Aris Koutsiamanis. It was implemented on the open-source Contiki OS and tested with the Cooja simulator. The DIO DAGMC NSA extension is implemented with a configurable number of parents from the parent set of a node to be reported.

(R)

```

(11)  (12)  (13)  (14)  (15)  (16)

(21)  (22)  (23)  (24)  (25)  (26)

(31)  (32)  (33)  (34)  (35)  (36)

(41)  (42)  (43)  (44)  (45)  (46)

(51)  (52)  (53)  (54)  (55)  (56)

```

(S)

Figure 4: Simulation Topology

The simulation setup is:

Topology: 32 nodes structured in a regular grid as shown in Figure 4. Node S (source) is the only data packet sender and sends data to node R (root). The parent set of each node (except R) is all the nodes in the immediately higher row, the immediately above 6 nodes. For example, each node in {51, 52, 53, 54, 55, 56} is connected to all of {41, 42, 43, 44, 45, 46}. Nodes 11, 12, 13, 14, 15, and 16 have a single upwards link to R.

MAC: TSCH with 1 retransmission

Platform: Cooja

Schedule: Static, 2 timeslots per link from each node to each parent in its parent set, 1 broadcast EB slot, 1 sender-based shared timeslot (for DIO and DIS) per node (total of 32).

Simulation lifecycle: Allow link formation for 100 seconds before starting to send data packets. Afterward, S sends data packets to R. The simulation terminates when 1000 packets have been sent by S.

Radio Links: Every 60 s, a new Packet Delivery Rate is randomly drawn for each link, with a uniform distribution spanning the 70% to 100% interval.

Traffic Pattern: CBR, S sends one non-fragmented UDP packet every 5 seconds to R.

PS extension size: 3 parents.

Routing Methods:

- * RPL: The default RPL non-PRE implementation in Contiki OS.
- * 2nd ETX: PRE with a parent selection method which picks as AP the 2nd best parent in the parent set based on ETX.
- * CA Strict: As described in Section 3.1.
- * CA Medium: As described in Section 3.2.

Simulation results:

Routing Method	Average Packet Delivery Rate (%)	Average Traversed Nodes/packet (#)	Average Duplications/packet (#)
RPL	82.70	5.56	7.02
2nd ETX	99.38	14.43	31.29
CA Strict	97.32	9.86	18.23
CA Medium	99.66	13.75	28.86

Table 1: Simulation results

Links:

- * Contiki OS DIO DAGMC NSA extension (draft-koutsiamanis-roll-nsa-extension branch) (<https://github.com/ariskou/contiki/tree/draft-koutsiamanis-roll-nsa-extension>)
- * Wireshark dissectors (for the optional PS TLV) - currently merged / in master (<https://code.wireshark.org/review/gitweb?p=wireshark.git;a=commit;h=e2f6ba229f45d8ccae2a6405e0ef41f1e61da138>)

Appendix B. Choosing an AP selection policy

The manner of choosing an AP selection policy is left to the implementation, for maximum flexibility.

For example, a different policy can be used per traffic type. The network configurator can choose the CA Relaxed policy to increase reliability (thus producing some flooding) for specific, extremely important, alert packets. On the other hand, all normal data traffic uses the CA Strict policy. Therefore, an exception is made just for the alert packets.

Another option would be to devise a new disjoint policy, where the paths are on purpose non-correlated, to increase path diversity and resilience against whole groups of nodes failing. The disadvantage may be increased jitter.

Finally, a network configurator may provide the CA policies with a preference order of Strict > Medium > Relaxed as a means of falling back to more flood-prone policies to maintain reliability.

Authors' Addresses

Remous-Aris Koutsiamanis (editor)
IMT Atlantique
Office B220
4 rue Alfred Kastler, CS 20722
44307 Nantes Cedex 3
France
Email: aris@ariskou.com

Georgios Papadopoulos
IMT Atlantique
Office B00 - 114A
2 Rue de la Chataigneraie
35510 Cesson-Sevigne - Rennes
France
Phone: +33 299 12 70 04
Email: georgios.papadopoulos@imt-atlantique.fr

Nicolas Montavont
IMT Atlantique
Office B00 - 106A
2 Rue de la Chataigneraie
35510 Cesson-Sevigne - Rennes
France

Phone: +33 299 12 70 23
Email: nicolas.montavont@imt-atlantique.fr

Pascal Thubert
Cisco Systems
Building D
45 Allee des Ormes - BP1200
3244 MOUGINS - Sophia Antipolis
France
Phone: +33 497 23 26 34
Email: pthubert@cisco.com