

ROLL Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 8, 2011

N. Dejean
Elster SAS
D. Barthel
France Telecom Orange
March 7, 2011

Selective DIS for RPL
draft-dejean-roll-selective-dis-00

Abstract

This document specifies DIS options that enrich the potential behavior of the Routing Protocol for Low Power and Lossy Networks (RPL) specified in [I-D.ietf-roll-rpl].

The goal is to enable new leaf nodes to quickly discover and attach to the routing structure, without having to wait for spontaneous DIO transmissions by neighbour routers and without causing them to reset their DIO timers.

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 8, 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. Leaf Node bit	4
3. DIS Options	5
3.1. Metric Container	5
3.2. Response Spreading	5
4. Example of use	6
5. IANA Considerations	9
5.1. DIS Flag Field	9
5.2. RPL Control Message Options	9
6. Security Considerations	9
7. Acknowledgements	9
8. References	9
8.1. Normative references	9
8.2. Informative references	10
Authors' Addresses	10

1. Introduction

This document makes use of the terminology defined in [I-D.ietf-roll-terminology].

Low power and Lossy Networks (LLNs) have specific routing characteristics compared with traditional wired or ad-hoc networks that have been spelled out in [RFC5548], [RFC5673], [RFC5826] and [RFC5867].

[I-D.ietf-roll-rpl] has specified the minimally viable core of mechanisms for a routing protocol, called Routing Protocol for Low Power and Lossy Networks (RPL), specifically designed for LLNs.

This document specifies DIS options that enrich the behavior of RPL and that were left out of [I-D.ietf-roll-rpl] in the interest of time.

The goal is to enable new leaf nodes to quickly discover and attach to the routing structure, without having to wait for spontaneous DIO transmissions by neighbour routers and without causing them to reset their DIO timers.

Indeed, with RPL as defined in [I-D.ietf-roll-rpl], a leaf node that wants to join an already deployed LLN is confronted with the following dilemma:

- o It can either wait for DIOs to be sent by neighbor routers. These transmissions may happen after a very long time, since the Trickle timers of the neighbor routers may have increased their period to a very large value, in order to save energy in a stable network. Furthermore, the transmission of a DIO packet by a neighbor router is not even guaranteed to happen during a Trickle timer period, since transmission suppression may happen (see [I-D.ietf-roll-trickle]).
- o Or it elects to proactively send a DIS (DODAG Information Solicitation). This DIS can only be sent in broadcast, since the new node does not know which router to ask for. Under the specification of [I-D.ietf-roll-rpl], all routers that receive a broadcast DIS packet will reset their Trickle timer. The time to their next spontaneous DIO transmission will indeed be dramatically shortened, which is desirable, although it will not prevent potential transmission suppression. But an undesired effect is that this will induce a large energy consumption in the network for two compounding reasons: first, all neighbour routers will respond, irrespective of their relevance to the new node, and second, each neighbor router will send frequent DIOs until its

Trickle timer relaxes to the maximum period, even though only the first DIO is useful.

None of the choices above matches the requirements of [RFC5548].

This document defines a way to broadcast a DIS message that includes selective options and a flag in order to query responses by neighbor routers such that:

- o responses are sent promptly, reducing the time the technician has to sit waiting at the customer premises to check the result of the joining process
- o responses are DIOs sent using unicast, reducing the overhearing energy cost in the router neighborhood when modern MAC technologies are used
- o each neighbor router only responds with a single DIO for each DIS, reducing the reception cost at the destination
- o the DIO is only sent if the neighbor router matches the criteria specified in the DIS selective options, reducing the reception, collision and overhearing energy costs

Admittedly, requesting an unknown population of neighbor routers to promptly send even a single DIO may be a cause for multiple collisions. This risk is mitigated by the use of good access contention methods at the link layer and by the wise use of the DIS options. However, both conditions are beyond the control of this specification. This document therefore specifies an optional collision mitigation mechanism of its own.

2. Leaf Node bit

In the format of the DIS base object, bit 0 of the Flag field is defined as the "Leaf Node" bit.




Figure 1: The DIS Base Object

A node that receives a DIS with the "Leaf Node" bit set MUST NOT reset its DIO Trickle timer, even if it matches the options carried by the DIS.

A node that receives a DIS message with the "Leaf Node" bit set and that matches the options carried in the DIS MUST reply with a unicast DIO, using the mechanism described in Section 3.2.

3. DIS Options

3.1. Metric Container

In addition to those already listed in [I-D.ietf-roll-rpl], the following option is declared valid for a DIS message:

0x02 Metric Container

A node that receives a DIS with a Metric Container option MUST ignore any Metric object in it, and MUST parse the Constraint objects in it, if any. The constraint values are compared to the values of the corresponding metrics known to the node. If both a Solicited Information option and a Metric Container option are present in a DIS message, they combine in a logical AND fashion, i.e. all predicates MUST match for the DIS to globally match.

If a Constraint objects carries a constraint for a metric the value of which is unknown to the node, it is RECOMMENDED that the node considers the constraint a match.

3.2. Response Spreading

With a wise use of the DIS options, our experience is that the population of responding routers is small enough for modern medium access techniques to efficiently resolve contention at the link layer. However, for those systems in which either above-mentioned postulate can't be met, an optional DIO response spreading mechanism is specified here.

A new RPL control message option is defined, called "Response

Spreading", with a recommended Type value of 0x0A (to be confirmed by IANA). Its format complies with the general format of RPL options, and is described in Figure 2.




Figure 2: The Response Spreading option

A node that responds to a broadcast DIS in observance of Section 2 MUST, if that DIS includes a Response Spreading option, wait for a time uniformly drawn in the interval $[0..2^{\text{SpreadingInterval}}]$, expressed in ms, before attempting to transmit its DIO. If the DIS does not include a Response Spreading option, the node is free to transmit the DIO as it otherwise would.

4. Example of use

A new leaf node that joins an established network runs an iterative algorithm by which it requests (using broadcast) network information from routers belonging to the desired network ID and which match some constraint values passed as parameters of the request. At each unsuccessful iteration, the requirements are relaxed, until one or several answers are received, or until the maximum number of iterations is reached. The answers from the routers can advantageously contain the values for other metrics than those by which the request was qualified, so that the router selection takes place based on more metrics.

The following example shows requests iterating on two constraint values (on Hop Count and Link Quality Level) and makes use of a third metric value (Node Energy) provided into the answers.

With Hop Count iterating through four different values (0-3) and Link Quality Level iterating through three possible values (2,4,6), a maximum of twelve DIS packets can be broadcast per joining node, in the following order:

- o Soliciting info from routers with max Hop Count 0 and max LQL 2
- o Soliciting info from routers with max Hop Count 0 and max LQL 4

- o Soliciting info from routers with max Hop Count 0 and max LQL 6
- o Soliciting info from routers with max Hop Count 1 and max LQL 2
- o Soliciting info from routers with max Hop Count 1 and max LQL 4
- o Soliciting info from routers with max Hop Count 1 and max LQL 6
- o Soliciting info from routers with max Hop Count 2 and max LQL 2
- o Soliciting info from routers with max Hop Count 2 and max LQL 4
- o Soliciting info from routers with max Hop Count 2 and max LQL 6
- o Soliciting info from routers with max Hop Count 3 and max LQL 2
- o Soliciting info from routers with max Hop Count 3 and max LQL 4
- o Soliciting info from routers with max Hop Count 3 and max LQL 6

Receiving any answer stops the iteration. Per our example, the new node then selects its parent router, based on the Node Energy and the Link Quality Level, according to the following algorithm:

- o Reject router(s) with asymmetric connectivity (LQL seen from new node does not match the constraint value issued in the DIS request)
- o Retain the router(s) that advertise the best Node Energy level
- o In case of equality, select the router that boasts the best Link Quality Level.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
155										0x00										Checksum																			
DIS BASE										Reserved										Solicited Information																			
L	Flags					Reserved					Type					Opt Length																							
1	0					0x00					7					19																							
Solicited Information																																							
RPLInstanceID										V I D					Flags					DODAG ID																			
=0x66										0 1 0					0					0x0000																			
Solicited Information																																							
DODAG ID																																							
0x00000000																																							
Solicited Information																																							
DODAG ID																																							
0x00000000																																							
Solicited Information																																							
DODAG ID																																							
0x00000000																																							
Solicited Information															MetricContainer																								
DODAG ID										Version Number					Type																								
0x0000										0x00					2																								
Metric Container																																							
Opt Length					Routing-MC-Type					Res Flags					P C O R A					Prec																			
12					3 (HC)					0					0 1 0 0 000					000																			
Metric Container																																							
Length (bytes)					Res					Flags					Hop Count					Routing-MC-Type																			
2					0					0					0					6																			
Metric Container																																							
Res Flags					P C O R A					Prec					Length (bytes)					Res																			
0					0 1 0 0 000					000					2					0x00																			
MetricContainer																																							
Val					Counter																																		
2					0																																		

Packet dump of DIS with Hop Count = 0, LQL <= 2

5. IANA Considerations

5.1. DIS Flag Field

IANA is requested to allocate bit 0 of the DIS Flag Field to become the "Leaf Node" bit, the functionality of which is described in Section 2 of this document.

Value	Meaning	Reference
0	Leaf Node	This document

5.2. RPL Control Message Options

IANA is requested to allocate a new code point in the "RPL Control Message Options" registry for the "Response Spreading" option, the behavior of which is described in Section 3.2.

Value	Meaning	Reference
0x0A	Response Spreading	This document

RPL Control Message Options

6. Security Considerations

7. Acknowledgements

8. References

8.1. Normative references

[I-D.ietf-roll-rpl]

Winter, T., Thubert, P., Brandt, A., Clausen, T., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., and J. Vasseur, "RPL: IPv6 Routing Protocol for Low power and Lossy Networks", draft-ietf-roll-rpl-18 (work in progress), February 2011.

[I-D.ietf-roll-trickle]

Levis, P., Clausen, T., Hui, J., Gnawali, O., and J. Ko, "The Trickle Algorithm", draft-ietf-roll-trickle-08 (work in progress), February 2011.

in progress), January 2011.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

8.2. Informative references

- [I-D.ietf-roll-terminology] Vasseur, J., "Terminology in Low power And Lossy Networks", draft-ietf-roll-terminology-04 (work in progress), September 2010.
- [RFC5548] Dohler, M., Watteyne, T., Winter, T., and D. Barthel, "Routing Requirements for Urban Low-Power and Lossy Networks", RFC 5548, May 2009.
- [RFC5673] Pister, K., Thubert, P., Dwars, S., and T. Phinney, "Industrial Routing Requirements in Low-Power and Lossy Networks", RFC 5673, October 2009.
- [RFC5826] Brandt, A., Buron, J., and G. Porcu, "Home Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5826, April 2010.
- [RFC5867] Martocci, J., De Mil, P., Riou, N., and W. Vermeylen, "Building Automation Routing Requirements in Low-Power and Lossy Networks", RFC 5867, June 2010.

Authors' Addresses

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

Email: nicolas.dejean@coronis.com

Dominique Barthel
France Telecom Orange
28 chemin du Vieux Chene, BP 98
Meylan, 38243
France

Email: dominique.barthel@orange-ftgroup.com